

PERSPECTIVE • OPEN ACCESS

Probe microscopy is all you need*

To cite this article: Sergei V Kalinin *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 023001

View the [article online](#) for updates and enhancements.

You may also like

- [Classical versus quantum models in machine learning: insights from a finance application](#)
Javier Alcazar, Vicente Leyton-Ortega and Alejandro Perdomo-Ortiz
- ['Flux+Mutability': a conditional generative approach to one-class classification and anomaly detection](#)
C Fanelli, J Giroux and Z Papandreou
- [Deep learning in electron microscopy](#)
Jeffrey M Ede



PERSPECTIVE





Probe microscopy is all you need*

OPEN ACCESS

RECEIVED
29 October 2022REVISED
26 January 2023ACCEPTED FOR PUBLICATION
13 April 2023PUBLISHED
22 May 2023

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.

Sergei V Kalinin^{1,**}, Rama Vasudevan^{2,**} , Yongtao Liu² , Ayana Ghosh³ , Kevin Roccapiore²
and Maxim Ziatdinov^{2,3,**} ¹ Department of Materials Science and Engineering, University of Tennessee, Knoxville, TN 37996, United States of America² Center for Nanophase Materials Sciences, Oak Ridge National Laboratory, Oak Ridge, TN 37831, United States of America³ Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, United States of America

** Authors to whom any correspondence should be addressed.

E-mail: sergei2@utk.edu, vasudevanrk@ornl.gov and ziatdinovma@ornl.gov**Keywords:** machine learning, microscopy, active learning, Bayesian, reinforcement learning**Abstract**

We pose that microscopy offers an ideal real-world experimental environment for the development and deployment of active Bayesian and reinforcement learning methods. Indeed, the tremendous progress achieved by machine learning (ML) and artificial intelligence over the last decade has been largely achieved via the utilization of static data sets, from the paradigmatic MNIST to the bespoke corpora of text and image data used to train large models such as GPT3, DALL·E and others. However, it is now recognized that continuous, minute improvements to state-of-the-art do not necessarily translate to advances in real-world applications. We argue that a promising pathway for the development of ML methods is via the route of domain-specific deployable algorithms in areas such as electron and scanning probe microscopy and chemical imaging. This will benefit both fundamental physical studies and serve as a test bed for more complex autonomous systems such as robotics and manufacturing. Favorable environment characteristics of scanning and electron microscopy include low risk, extensive availability of domain-specific priors and rewards, relatively small effects of exogenous variables, and often the presence of both upstream first principles as well as downstream learnable physical models for both statics and dynamics. Recent developments in programmable interfaces, edge computing, and access to application programming interfaces (APIs) facilitating microscope control, all render the deployment of ML codes on operational microscopes straightforward. We discuss these considerations and hope that these arguments will lead to create novel set of development targets for the ML community by accelerating both real world ML applications and scientific progress.

The last decade has seen the meteoric growth of methods and applications of deep learning (DL). Starting with the seminal paper by Heravi *et al* [1] that introduced DL for image recognition tasks trained with the ImageNet database by Deng [2], the field has developed almost exponentially, as evidenced by both numbers of papers, as well as recognition of the field in everyday society. The early overviews [3, 4] have laid the landscape for the applications of DL in a variety of areas, from medicine to computer vision to robotics. Over the last 2–3 years, we have witnessed an accelerating trend of adoption of DL in more experimental sciences, including chemical synthesis [5], materials science and condensed matter physics [6, 7], as well as multiple modalities of chemical and physical imaging [8].

* Notice: This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

The rapid development in DL architectures for classical neural network architectures resulted in new models including different forms of convolutional networks [9], introduction of graph networks (e.g. crystal graph convolutions) [10], and transformers [11], by now ubiquitous for natural language and increasingly, vision applications. It has also stimulated the development of generative models such as variational autoencoders (VAEs) [12–14], and generative adversarial networks (GANs) [15] and their variants that enable a broad range of generative and style transfer applications [16–19]. Finally, advancements in deep reinforcement learning gave rise to such examples as muZero [20], AlphaZero [20, 21], OpenAI5 [22], and real-world applications such as in control over a robotic hand [23], as well as navigation of stratospheric balloons [24].

It is important to note that this progress has been based on multiple algorithmic and technical developments over the preceding decades. Neural networks were developed as an idea by McCulloch and Pitts in 1943 [25]. Advances in computing would facilitate the invention of the first neural network *in-silico*, termed the ADALINE network [26], which was used to filter audio signals. Multi-layer neural networks, studied first by Rosenblatt [27], were explored in the 1970s [28], but they were not significantly utilized until the backpropagation method was developed in 1986 by Rumelhart [29] to enable efficient training. Reinforcement learning has been known for similarly long periods of time, going back to the works of Widrow in the late 60s [30], with other pioneering work by Sutton [31]. The q -learning algorithm was developed in 1989 by Watkins [32]. The long short-term memory network was introduced by Schmithuber in 1997 [33], and recurrent neural networks were studied for a decade before that seminal work. However, it is only over the last decade that the confluence of available computation capabilities [34] and large data sets, emerging in the context of internet searches, social media, and open knowledge projects, that these methods became practically useful across multiple domains and applications.

In parallel with the development of algorithmic tools, the community has developed a range of static data sets, mainly driven by the need to define and benchmark the machine learning (ML) methods. In classical image analysis these include the paradigmatic MNIST hand written digits [35], and multiple CIFAR data sets [36]. Similar data sets have been developed in domain areas, including theoretical quantum science, e.g. with the NTangled dataset [37], several examples in climate science [38–40], as well as in materials science, for example the SuperMat dataset for superconducting materials [41], or that of multi-principal element alloys [42]. Note that there is also a large number of datasets of simulated data in materials science, many of which are discussed in recent reviews [43].

Much progress over this last decade has been fueled by the availability of these well-established datasets enabling benchmarking. While remaining a topic of considerable controversy, it is being argued currently that developing of the large ‘foundation’ models such as GPT and DALLE may be almost exhausted [44]. Indeed, these models clearly show outstanding performance for in-distribution tasks. However, they very often fail to generalize or yield a well-defined answer for a specific question (as anyone who played with DALLE or Craiyon have discovered). Unsurprisingly, the first applications of DALLE in science is cover art. Complementing these, somewhat less visible are problems of extreme training in costs that make these models available only to large research organizations (and very often trained models are proprietary). For example, the compute time for GPT-3 was 3650 petaflop-days [45], which is about 34 d of continuous use of 1024 T A100 graphic processing units (GPUs). The energy for this training is about 190 000 kWh, which is about what 19 US homes use on average in an entire year [46]. Similarly, emerging is the issue of data exhaustion, i.e. models can be trained with full Wikipedia only once.

Complementary to it are by now multiple cautionary stories where DL models trained on static data sets are limited, and fail to generalize. Multiple examples of this problem abound, for example in self-driving cars, radiology and biological imaging [47–52]. It was the experience of the authors of this opinion piece that deep convolutional neural networks (DCNNs) trained on theoretical (i.e. simulated) data requires tuning for specific microscope settings [53], but will fail to generalize for a different set of microscope parameters. While solutions based on more complex contrastive losses such as Siamese networks [54] and Barlow twins [55], ensembling [56], and invariant risk minimization [57] are being continuously developed, this problem is far from resolved.

A number of the leading scientists now argue that artificial intelligence (AI) needs to move away from ever-increasing dataset and network sizes, to essentially ‘do more with less’ [58] by using small data sets and physical priors. Similar arguments have been made in the physical science field [59, 60], typically in cognizance of rich systems of prior knowledge and inferential biases and extremely small experimental budgets. This convergence between ML and physical fields is, in our opinion, a significant aspect of the last three years.

Traditionally, it is believed that ML methods emerge at the interface between classical computer science and established applications such as robotics, automated vehicles, and computer vision, are developed within these domains, and then become accepted in domain communities. However, for the domain experts it is

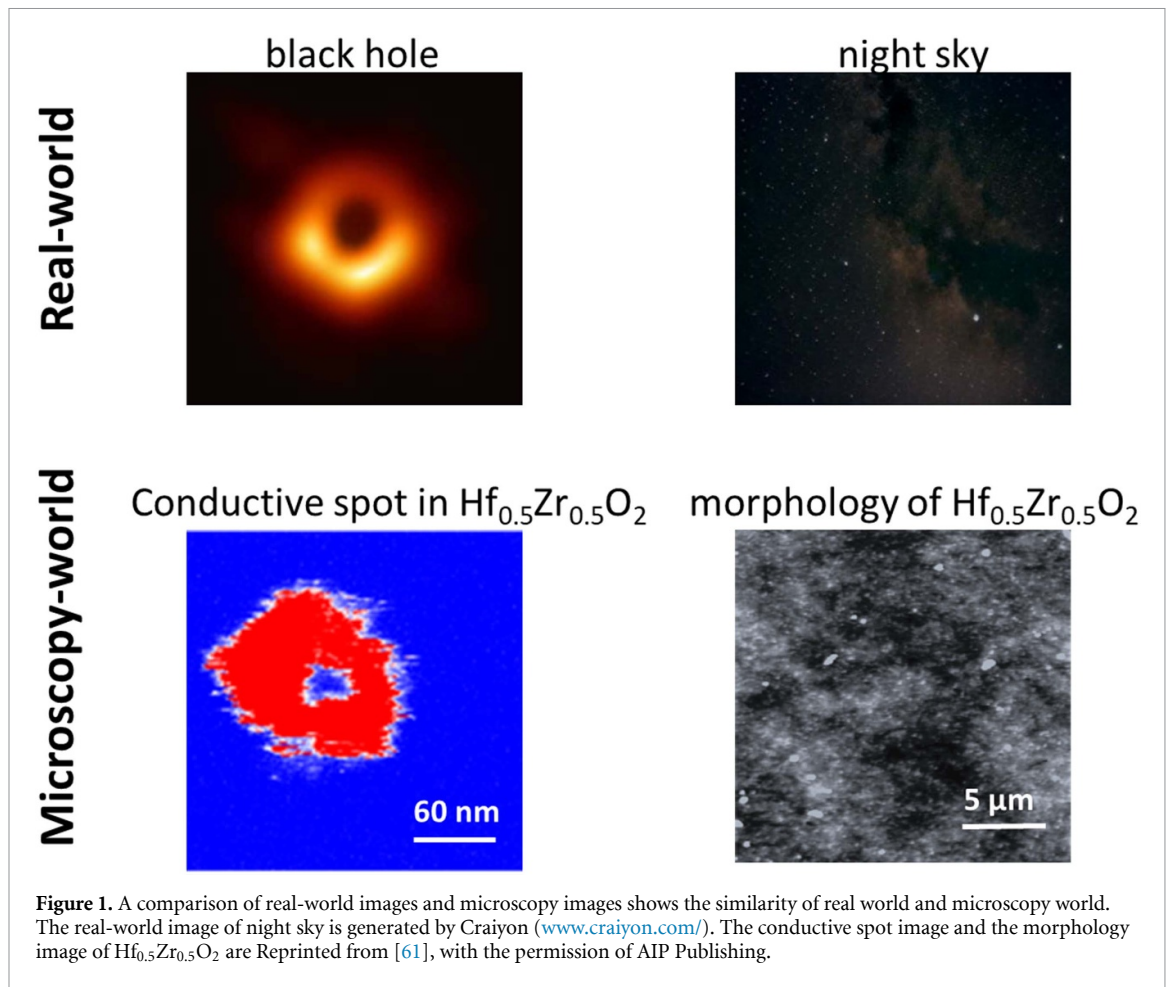


Figure 1. A comparison of real-world images and microscopy images shows the similarity of real world and microscopy world. The real-world image of night sky is generated by Craiyon (www.craiyon.com/). The conductive spot image and the morphology image of $\text{Hf}_{0.5}\text{Zr}_{0.5}\text{O}_2$ are Reprinted from [61], with the permission of AIP Publishing.

Table 1. Microscopy is all you need (for now).

Autonomous driving	Electron microscopy	Scanning probe microscopy
Camera glitch	Cosmic ray spark, electron gun change	Tip changes
Weird weather	Unexpected surface contaminations	Surface contamination
Camera out of focus	Microscope focus	Setpoint setting/sample drift
Optimize motor	Tune the microscope	Tune feedback
Conduct actions (turns, navigation)	Navigate to different parts of sample	Navigate to different parts of sample
Simulated drives for data collection	Simulated images from scattering codes	Simulated images from phase-field models or density functional theory
Collision	E-beam induced sample damage	Tip crashes into a surface
N/A	Modify material with electron beam	Modify material with tip pressure or bias

immediately clear that adoption of the problems and challenges in the ML community often follows an ‘inverse’ logic, with inherently complex and high-risk problems accruing most of the attention and investment, whereas relatively simple and well-defined problems are ignored or explored only in the domain context. At the same time, many domain problems offer highly simplified toy problems for ML. For example, in figure 1 and table 1, we establish the similarity between the challenges in the context of AI for self-driving vehicles and modern microscopies.

In this opinion, we argue that a promising new venue for the development of active learning methods capable of operating in real world situations are electron and scanning probe microscopy (SPM), both for intrinsic discovery value, and as simplified models for more complex systems. These methods are relatively low risk (compared to automated driving), are usually fairly closed systems minimizing but not excluding exogenous effects (unlike e.g. stock-market), come with a multitude of domain specific reward functions, allow for multiple levels at which prior knowledge can be incorporated, rely on systems with well-defined but often unknown or partially known physical models, and allow for interventional and counterfactual studies. Most importantly, over the last few years, Python application programming interfaces (APIs) and remote

operation has become broadly available, allowing for the direct (via edge computing) and remote deployment of the codes on operational machines and their digital twins. Below, we expound on these considerations and (following a well-established ML meme) argue that electron and SPM is what the ML community needs (for now).

1. What are electron and SPM

When writing this opinion, we feel it is prudent to briefly introduce scanning transmission electron microscopy (STEM) and SPM to an ML community. STEM is an imaging technique based on the electron beam, as illustrated in figure 2(a) [62]. The beam is emitted by the electron gun and subsequently focused and monochromated via a bespoke collection of electron lenses. Effectively the beam is compressed to sub-atomic width and formed by electrons with a very narrow energy distribution. The STEM sample is a very thin membrane, typically below 50 nm thickness, i.e. ~ 1000 atoms. These samples can be fabricated via a variety of methods ranging from the native formation via crushing or delamination to highly sophisticated techniques such as focused ion beam milling. Generally, the infrastructure for this sample fabrication is by now well established, and is available in academic, national labs, and industrial centers, and we will not expand further on it here.

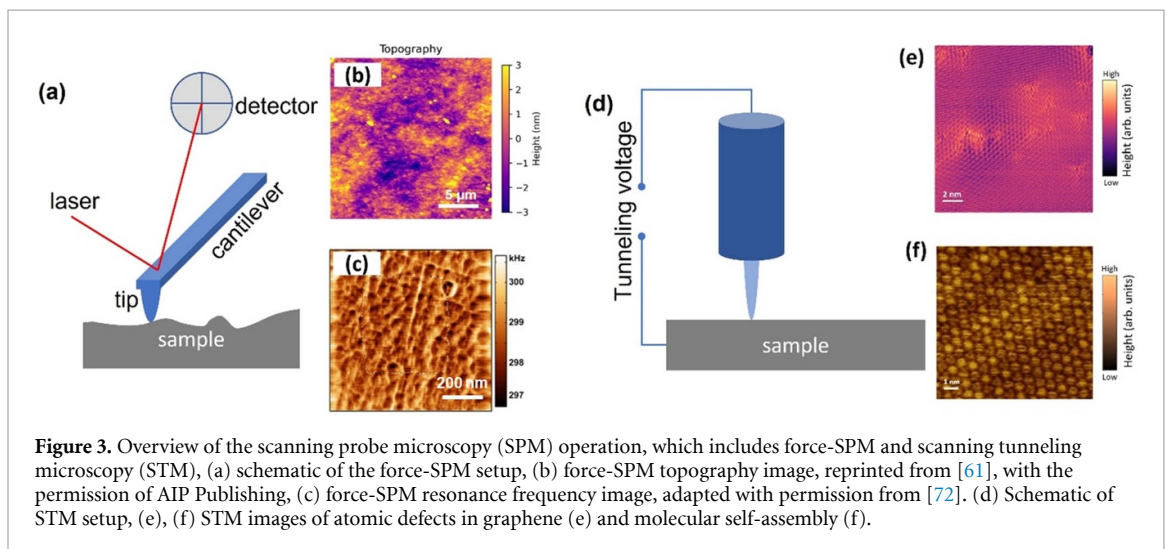
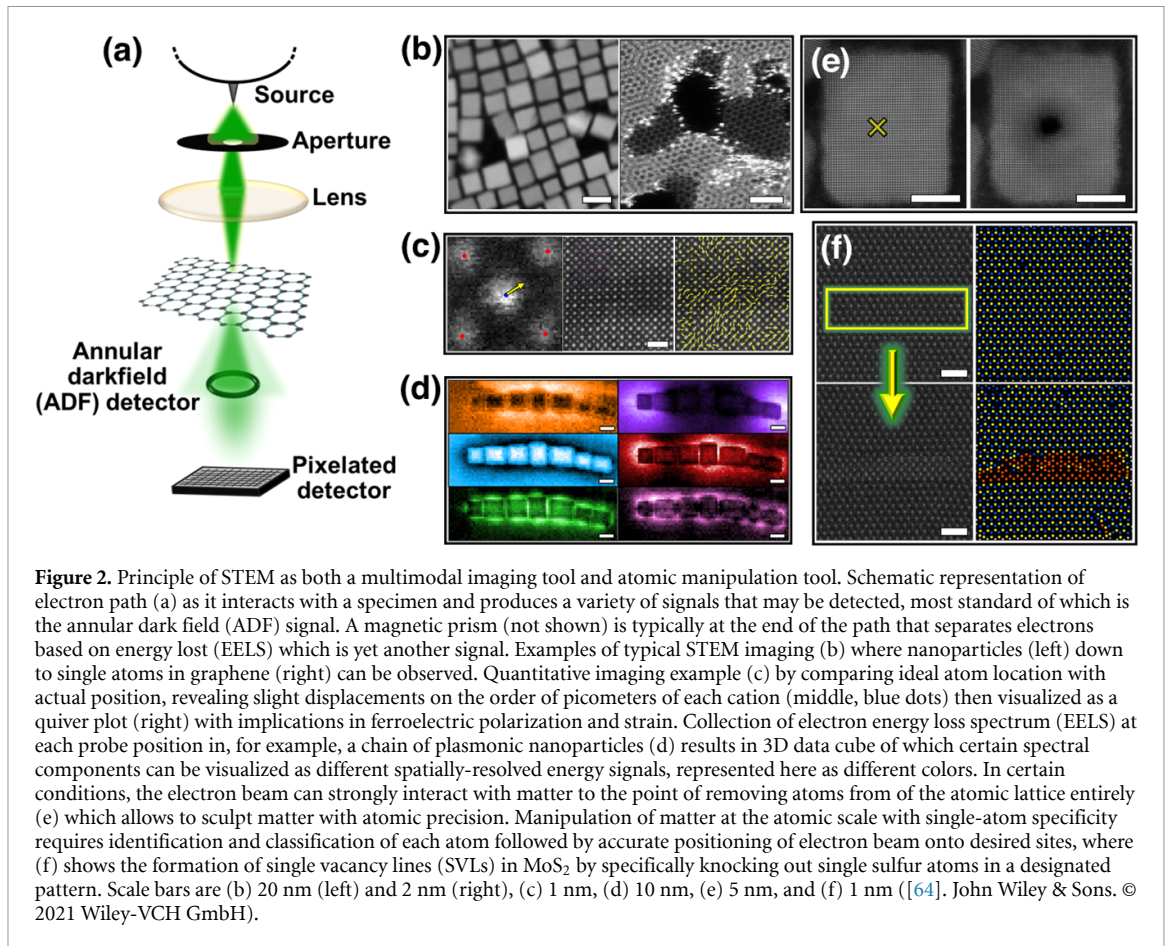
Upon transmission through the sample, the electrons are collected via bright field (direct transmission) and dark field (electrons that were partially deflected) detectors. The signal intensity thus provides the information on the density of the solid. Perhaps the simplest analogy for STEM image formation mechanisms is shining the flashlight onto the construct and observing the shadows that transmitted light makes on the wall. Despite apparent simplicity, STEM allow direct visualization of atomic structures, with often spectacular images of atomic structures of metals, oxides, and semiconductors being available. In certain cases, the STEM images can be quantified, providing direct information on atomic spacing and bond length [63]—via direct images!

In addition to imaging, STEM allows a number of other interesting possibilities. The energy distribution of transmitted electrons can be measured (much like optical spectra) in a technique called electron energy loss spectroscopy (EELS). These EELS spectra can also be localized on the level of single atoms [65] and provide information on chemical and orbital states of individual atoms, energy excitations and quasiparticles, etc. The STEM is highly relevant to applications such as quantum computing (beam induced photon emission, quasiparticles and other excitations). Most intriguingly, the electron beam can be used for atomic assembly [66, 67], to remove or move single atoms [68, 69] and assemble atomic structures [67, 70]. For the time being, with few exceptions [71] the atomic manipulation was enabled with purely human operator control.

The second type of microscopy we discuss here is the SPM. SPM is based on the interaction of the sharp mechanical probe, or tip, with the surface. In force-based SPMs, the probe is fabricated on the end of a silicon cantilever and the interaction between the probe and the surface leads to the deflection of the cantilever (figure 3). This deflection is then detected via either an interferometer or a laser beam deflection system [73]. Despite simplicity, the sensitivity of this method is such that the forces on the level of interatomic ones can be detected, hence the name atomic force microscopy [74]. The simplest but apt analogy to SPM is touching the surface with the finger (with eyes closed).

Depending on the specific configuration and operation mode, force-based SPMs can be made sensitive to a broad range of forces including magnetic, electric, and Van-der-Waals, allowing visualization of phenomena such as magnetic domains [75, 76], electronic and ionic transport [77–79], mechanical properties of the surface [80], ferroelectric domains [81], and many others. SPM can be implemented on active devices, exploring phenomena induced by lateral stimuli [78, 79, 82, 83]. The SPM tip can be further used to manipulate matter, from local nanooxidation and lithography [84] and writing ferroelectric domains [85] to mechanical manipulation, moving particle and nanowires around [86–88]. In fact, multiple versions of haptic interfaces for SPM have been developed over time providing a human interface to the nanoworld [89–91].

A special (and first invented) class of SPMs is scanning tunneling microscopy (STM) [92–94]. STM uses simple metallic probe and detects the tunneling current between the probe and conductive surface. While generally requiring a ultra-high vacuum environment for operation, STM is capable of visualizing surfaces on the atomic level and manipulating matter atom by atom. In fact, the letters ‘I’, ‘B’, and ‘M’ written in xenon atoms on the copper surface by Don Eigler at IBM [95] was the initial demonstration of atomic manipulation, and gave impetus to the nanotechnology revolution. Now, manipulation of phosphorus atoms on silicon surfaces is broadly considered as one of the possible pathways to solid-state qubits [96, 97] for quantum computing.



Common for both SPM and STEM is the basic principle of a mechanical (SPM, STM) or electron beam probe that can be scanned across the material surface and detect the local state (imaging mode), perform more detailed measurements (spectroscopy), or modify the material state. We note here that this is often a differentiating factor from standard optical microscopy, where the amount of functional property information that can be acquired is not as broad and ability to perturb the sample is generally limited (except via external means). The probe parameters and trajectory can be controlled by the microscope electronics. In this manner, we consider these to be systems that allow a limited set of operations and controls on the local level and provide a convenient model system for deployment of ML, active learning and reinforcement learning algorithms. Below, we discuss these considerations further.

2. Why microscopy for ML

Electron and SPM are, by now, well-established fields with broad applications within physics, chemistry, and materials science. Below, we discuss aspects of these field that, in our opinion, makes them ideal model systems for active learning methods and hence the ML community. In all cases, we introduce domain-specific considerations and illustrate the connection to the ML.

2.1. Language and control

The vast majority of STEM and SPM studies have been performed with commercial tools, often representing black box systems with fairly narrow user interfaces and (for STEMs) service contracts. While customizable controllers have been available (e.g. *Nanonis*) they used to represent a small part of the field. Correspondingly, deployment of custom developed spectroscopies and ML tools required teams comprising microscopy, physics, and electrical engineering expertise and have been the exception rather than the rule [98–101].

The situation started to change rapidly over the last several years, when a number of microscope manufacturers started to introduce Python programmable interfaces for commercial tools. A partial list includes SWIFT for NION microscopes [102], PyJEM for JEOL systems [103] and even more open source efforts in optical microscopy that are significantly hardware agnostic, such as pycroManager [104]. For SPM, this includes Python interfaces e.g. the NanoSurf atomic force microscopes [105], and home-grown efforts at multiple research centers worldwide [106, 107] for control of low-level microscopy functions on a range of instruments, all via python commands [108]. Correspondingly, the operator has significant programmable control over the microscope operation, and thus correspondingly, access to data streams. This makes it significantly easier to deploy ML codes on these data streams to affect microscope operations, thus greatly lowering the activation barrier for deployment.

It should be noted that one of the advantages of the ability to control the probe position externally is that it enables implanting fast scanning approaches, such as those using spiral scans or compressive sensing techniques, thus enabling significant speedups in image acquisition (on the order of ~ 1 – 2 Hz in SPM for example, and faster for electron microscopy) [109–111].

Additionally, the ability to ‘script’ entire experiments and run them on multiple instruments is a significant shift not only for the types of experiments that are then possible (due to automation), but also in that it changes the user paradigm to one of remote operation, in tune with the recently developed ‘cloud labs’ for chemical synthesis [112, 113]. This parallels the deployment of online quantum computers that allow users to run their own codes by simply uploading them. Remotely operated and automated microscopy instruments offer the same tantalizing possibility, to test not only new experimental workflows, but to benchmark ML algorithms across a range of microscopy platforms.

The last decade has also seen significant developments in terms of data models and associated infrastructure for the microscopy community. Although there exist many formats, on almost all occasions typically there exist ways to convert between multiple formats and read data and associated meta-data into python, regardless of which instrument the data was acquired on. For SPM, virtually every vendor produces their own data format, but programs such as WsXM [114] and Gwyddion [115] have long been capable of reading of data from a wide variety of SPM formats. Tools such as pySPM [116] and Pycroscopy SciFiReaders [117] have significant collections of ‘translators’, which build on community tools to convert data into standardized data models, such as the spectral imaging dataset [118]. Regardless of the choice of final format, the codebase exists to convert most existing datasets to a common one for generating curated ML datasets.

A similar situation exists in the electron microscopy, with translators existing for virtually all of the commonly used microscopes, with packages such as HyperSpy [119] and openNCEM [120] providing the necessary bridges to ensure interoperability. In contrast with most scanning probe work, it is interesting to note that the idea of parallel computation and workflows is more established within the TEM community through efforts such as LiberTEM [121], primarily due to the nature of the reconstruction algorithms that are compute intensive. Note however, that this is changing rapidly within SPM in the past two years, and workflow solutions more specific to SPM can be expected to follow suit in due course.

Finally, we note that a significant portion of development on microscopy platforms occurs at scientific user facilities, in both the United States and across other centers around the globe [122]. This dynamic ensures continuous development that is not limited to groups led by individual principal investigators (PIs) at universities, and furthermore, encourages a more broad-based effort, given the needs of the user center are to ensure new techniques and analysis methods are available to the user community with minimal barriers. However, it is still incumbent on the whole microscopy community to engage more deeply with data sharing

and curation [100]. Overall, we are not yet at the stage when the codes from GitHub can be deployed to run on selected microscope as a ‘physical RL gym’ from the ML side. However, from the domain side it is already possible to run specific code as a plug-in to the microscope.

2.2. Risk

One significant consideration in real world applications of ML is that of risk. Perhaps the most recognized case for it is autonomous driving, where each accident is instantly popularized. However, this is also well recognized in areas such as ML fairness and explainability, with the algorithm trained on historical data often making predictions and recommendations inconsistent with current social norms [123].

Comparatively, for microscopy the risk is relatively small. It is well understood that scientific research is prone to errors, and knowledge emerges as the result of multiple trial and error cycles, often including community feedback. Given this established paradigm, any results derived from the automated workflows will be verified by conventional methods for an extended time, and the impact of potentially incorrect findings can be expected to be comparable to that derived during normal research processes.

The second component is the risk to the instrument. While crashed kernels in the notebook can be restarted, the crash in a microscope often necessitates the change of the probe. For ambient SPMs this is of the order of ten minutes and process can be automated. For UHV systems it can be several days (and hence STM community is highly conservative). Beyond the probe, the damage for the mechanical parts of the systems may be considerable. Practically however, these risks can (and are) heavily mitigated by constraints in allowable operations. Following on the concepts and ideas from the operation of cloud compute infrastructure and automated labs, the allowed actions in automated microscopy can be chosen to be within safety tolerances to ensure equipment safety. Similarly, de-risking is possible via the introduction of the digital twins, meaning the *in-silico* models of the microscope that either emulate operation via ground-truth data, or allow modeling of the mechanical and electronic equivalent of the microscopes [124].

Note that the advantage of the microscope as a model system for active learning is the low risk combined with the large variability of microscope and sample environments, ideal for developing ML solutions stable with respect to large out of distribution (OOD) effects. Once developed, these can be potentially transferred to higher-risk problems, e.g. robotics or autonomous drones.

2.3. Exogenous variables and causal models

A very important consideration for any real-world ML algorithm is the number and character of exogenous variables, i.e. external (to the system) stochastic parameters. For example, the success of AlphaGo [21, 125] or any other reinforcement learning algorithm for computer games is predicated on known rules and lack of external variables (e.g. the number of squares on the board will never change). In other words, while the game contains the element of randomness, the general rules are constant, and the algorithm builds the approximations via the Q-tables or continuous policies.

Comparatively, the self-driving car or stock-trading algorithm has to deal with influences outside of the system—from pedestrians coming from around the corner to political events. Correspondingly, long-term prediction or even models of the environment for long time are difficult to construct. Clearly, the absence of exogenous variables is an oversimplification of the real-world applications (the world is not a chess game), while too many limits how far we can predict (it is not a fully random process either).

From this perspective, microscopy offers an intriguing range of possibilities from almost deterministic to highly stochastic systems and from full to partial knowledge of the system, all in the presence of strong physical priors. For example, STEM imaging of bulk materials visualizes the atomic columns averaged over the beam direction. Imaging with focal and tilt series, while more complex, provide information on the structural variability in the beam direction [126–130]. Atomically resolved images of 2D materials such as graphene allow all static atoms to be visualized in the image plane, but the observed phenomena may also be affected by the presence of (invisible) moving atoms that can participate in chemical reactions and atomic reconfigurations. At the same time, at this length scale the objects that we observe in the microscope can only be atoms, and atoms have a well-defined geometry. Correspondingly, ML analysis of such data allows for very strong inferential biases.

For SPM microscopy these considerations are more complex and can be less well defined depending on the specific material system. For example, UHV STM relies on atomically clean surfaces to minimize unwanted noises and chemical contamination (i.e. each stray atom will be seen as a ‘mound’ on surface). Techniques such as piezoresponse force microscopy [131] allow almost complete decoupling of information on different aspects of materials functionality, e.g. topography, elastic properties, and ferroelectric domain structures. At the same time, magnetic force microscopy [76] will be sensitive to both magnetic and electric interactions, whereas phase imaging mixes all interactions including magnetic, electrostatic, adhesion, and elastic into one information channel.

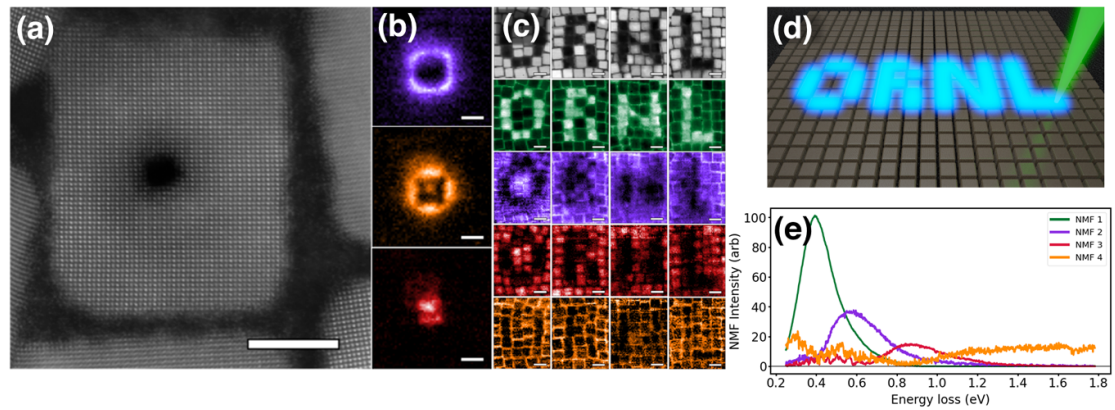


Figure 4. Electron beam sculpting of matter and its connection to its nanophotonic properties. Specific energetic features, represented by different colors, can be placed by the electron beam by sculpting nanoparticles into geometric shapes that support a desired nanophotonic property. Hyperspectral EELS image is acquired then decomposed into four dominant spectral components using non-negative matrix factorization (NMF) shown in (e), with corresponding abundance maps in (c). Scale bar in (a) 5 nm, (b) 10 nm, and (c) 20 nm. [64]. John Wiley & Sons. © 2021 Wiley-VCH GmbH.

Yet, compared to largely uncontrollable interactions in the real world, microscopy offers an opportunity to explore correlative and causal links in technologically relevant systems under (mostly) controlled environments, starting from relatively simple systems with largely known physics (e.g. one-element-based 2D materials) and progressing towards more complex systems demonstrating yet-to-be understood physical behaviors (e.g. cuprate superconductors).

For example, SPM and spectroscopy allow collecting data on both atomic positions and physical behaviors of interest (encoded in spectra) over the same sample region [132, 133]. The local atomic distortions and classical and quantum phenomena are assumed to be correlated and their relationship is considered to be parsimonious, that is, ultimately explained by a small number of (latent) mechanisms. A conventional approach to the analysis of such data is to apply a blind unmixing to the hyperspectral data and do a linear correlation analysis of the abundance maps of the derived sources with structural descriptors at each unit cell [134]. It would be interesting instead to utilize deep latent variable models trained in end-to-end fashion (with structural images as input and spectroscopic curves, full or scalarized, as outputs) to disentangle different physical mechanisms and establish how changes in specific geometric features influence a property of interest (for example, local T_c). Furthermore, in complex systems with competing interactions and chemical disorder, it will be critical to distinguish between correlative and casual links, offering a playground for a newly emerging class of casual representation learning techniques.

2.4. Interventions and counterfactuals

Electron and scanning probe microscopies are often perceived as a purely observational areas, much like astronomy. However, familiarity with the field suggest that this is generally not the case—both probes and beams can modify the surface, but these behaviors are often perceived as deleterious. For example, while a high energy electron beam can easily modify the sample, this is traditionally perceived as beam damage. In fact, much of the evolution of modern STEMs, first towards high voltages and then towards low voltage aberration corrected machines was driven by the need to balance resolution and beam damage. Similarly, in SPM the damage to the surface and especially the probe is generally unwanted.

However, for many materials classes the evolution of both SPM and electron microscopy techniques has facilitated control over the modifications inflicted on the sample. Atomic manipulation in STM [95] and STEM [68, 70, 135] is one example of this. Here, we create atomic configurations that are absent in the original material, and explore their functionality via imaging and spectroscopy (figure 4(b)). Similarly, in piezoresponse force microscopy we can create new domain structures and explore their polarization dynamics [81, 136].

From a ML perspective, this opens a pathway for exploring interventional and counterfactual questions implemented on an experimental platform. For example, a system of plasmonic nanoparticles contains multiple local geometries that can be extracted from the experimental data and correlated with the associated plasmonic spectra using im2spec [137, 138] or parallel VAE approaches [139, 140]. These methods allow us to predict the plasmonic properties for the extant geometries. However, the electron beam can be used to modify the system in a broad variety of scenarios, including those that remain in distribution compared to original data (e.g. remove the particle and create a hole similar to the ones that already exist), or create

geometries that were not present in the original data sets (e.g. hole within the particle). We note here that the ability to repeatedly modify and probe the system should also enable a playground for ‘concept learning’ for automated agents. Combined with the reasonably well-understood physics of these systems, this opens a very broad set of opportunities to implement causal active learning methods and explore interventional and counterfactual strategies.

2.5. Prior knowledge and active learning

Most ML methods are correlative and generally take only limited advantage of the prior knowledge, typically implicitly in the form of output or the loss function. As such, they capture the structure of the target distribution, but generally impose no limitations on the properties of the object. However, physical sciences take a very different approach where past knowledge plays the preponderant role, and the experiment is often formulated to falsify the specific hypothesis formulated based on prior domain specific and general knowledge.

Recently, significant progress has been achieved in developing physics-informed neural networks [141]. In these, the network structure can be such that outputs directly satisfy certain invariances or equations of motion. Alternatively, the known factors of variability such as rotations, translations, or shear can be incorporated as defined latent variables in architectures such as VAEs [142–144].

From the perspective of physics-informed ML, STEM and SPM offer a broad range of opportunities both on the instrument and materials side. From the instrument perspective, the SPM cantilever oscillations closely follow Euler beam equations, providing a strong inferential bias. In electron microscopy, the image formation mechanisms are complex, but the physics of the multiple scattering can be naturally incorporated into the neural network structure. An even broader range of opportunities emerge on the materials side, where materials systems chosen for detailed study usually come with known composition, preparation history, and hence well-defined physical priors in structure and range of possible phenomena.

The important aspect for ML applications is that the experimental setup in SPM is often selected to narrow down the scope of relevant knowledge. For example, when exploring ferroelectric domains [145], the relevant information with a good degree of approximation is the domain structure tied to the corresponding free energy functional and the preparation method. At the same time, a second level of parameters will be the surface preparation and surface chemistry conditions [146–148]. Often the potential contributions of certain level of description can be estimated from straightforward physics-based estimates, with the remaining variables contributing to exogenous factors with some known level of significance.

2.6. Physical models

The characteristic aspect of many physical systems is the simplicity of the laws that govern their behavior, as exemplified by the fundamental gravitational laws, Maxwell equations, or quantum mechanics. This local simplicity however gives rise to very complex behaviors of real systems, from real material properties to the emergent behaviors in systems with collective interactions [149]. The important aspects of these laws is that generally they are expected to be universal, e.g. the force fields between two carbon atoms do not depend on where in the universe these atoms are. Correspondingly, it can be argued the high-resolution imaging studies can be aimed at discovery and identification of these laws.

The significant progress toward these applications of ML has been accomplished in astronomy. It has been demonstrated by Lemos *et al* that the symbolic laws of gravitation can be derived from observations of planetary motion [150]. Similarly, Udrescu and Tegmark have demonstrated that equation of motions can be derived from video data [151].

From this perspective, microscopy data opens virtually unlimited opportunities to explore physical models. Provided only the position of atoms, it is possible to extract a number material properties. For example, built-in polarization fields and ferroelectric domains can be identified by—i.e. the electron beam in the STEM is sensitive enough to resolve picometer displacements that have origins in polarization [153]. Figure 5(a) demonstrates how the determined atomic coordinates in a ferroelectric superlattice are used to compute polarization fields—simply, the deviation of a cation (bright atom) from its ideal lattice position gives the polarization vector at that site. Similarly, it is known that strain can in turn have an impact on electronic properties [154], but it turns out only a small fraction of atoms deviate from their nominal lattice location, for example, at the interface of a core–shell nanoparticle. Hence, we can begin to understand the underlying reason why a material exhibits a certain response—it may be intimately tied to a very localized strain or built-in polarization field.

Another subtle feature that can be detected via atomic coordinates is the presence of defects, which can exist in the form of vacancies, impurity atoms, or practically anything that is not part of the pristine crystal. Many times, these can be extremely difficult for the human operator to notice (especially in the case of a

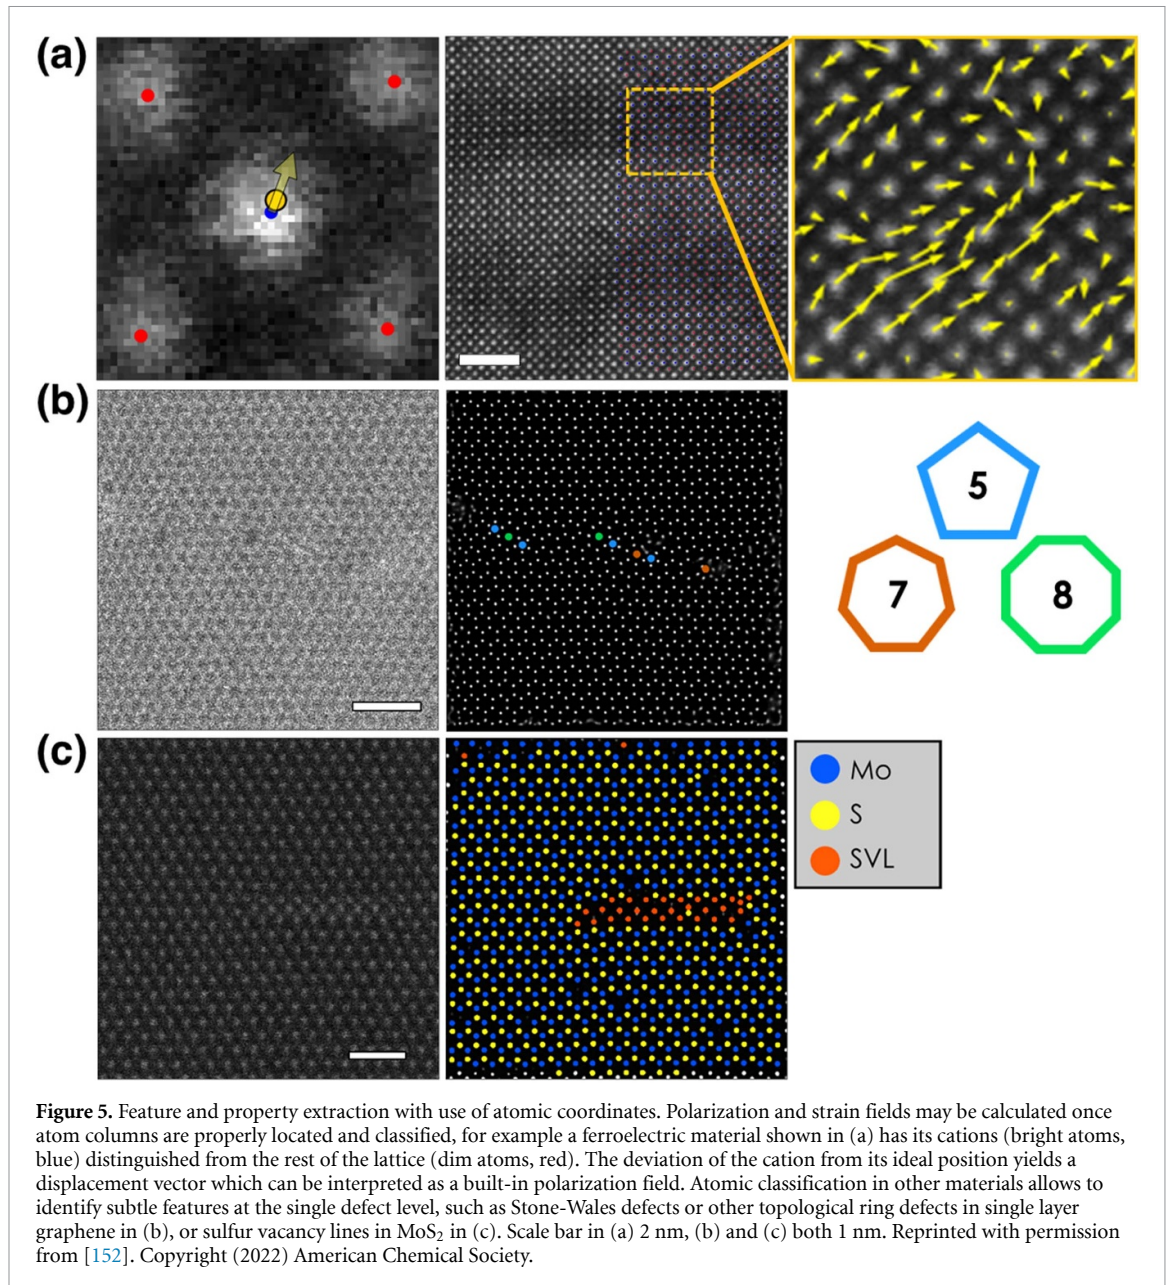
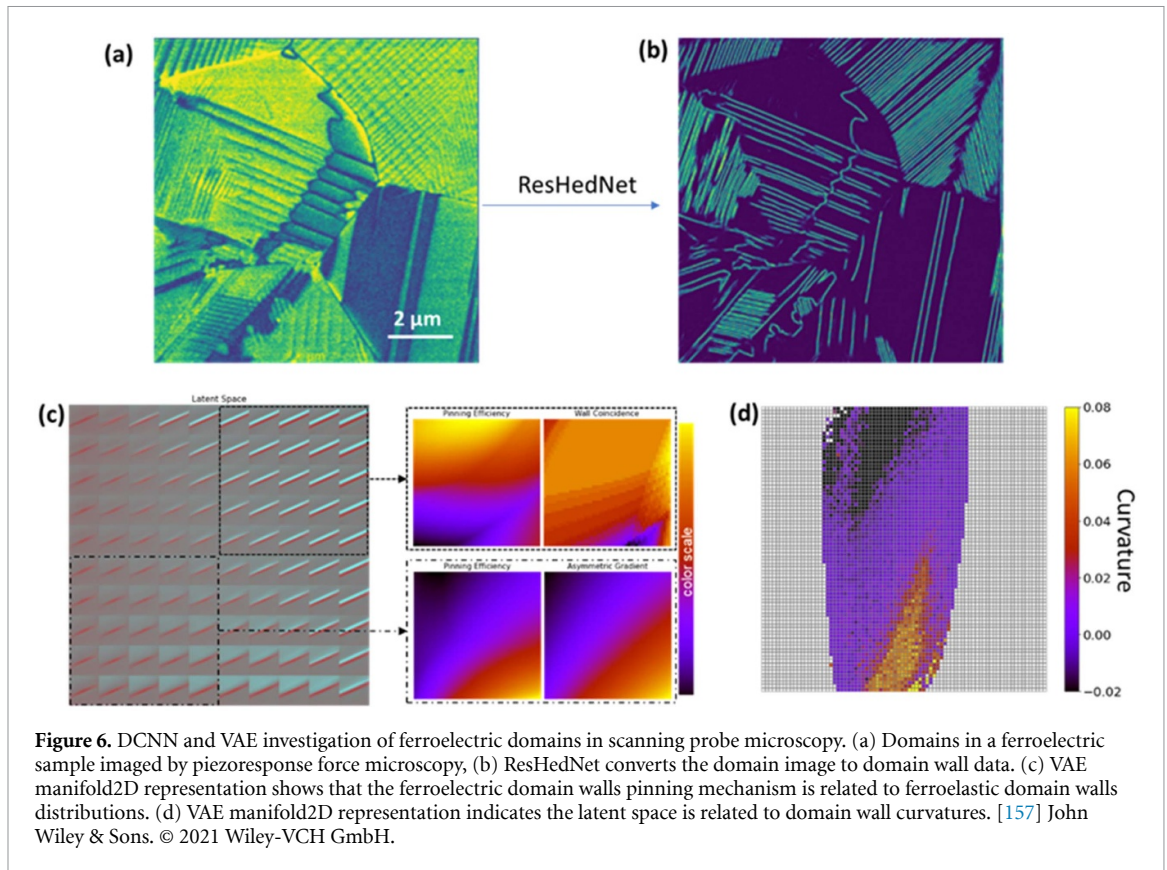


Figure 5. Feature and property extraction with use of atomic coordinates. Polarization and strain fields may be calculated once atom columns are properly located and classified, for example a ferroelectric material shown in (a) has its cations (bright atoms, blue) distinguished from the rest of the lattice (dim atoms, red). The deviation of the cation from its ideal position yields a displacement vector which can be interpreted as a built-in polarization field. Atomic classification in other materials allows to identify subtle features at the single defect level, such as Stone-Wales defects or other topological ring defects in single layer graphene in (b), or sulfur vacancy lines in MoS₂ in (c). Scale bar in (a) 2 nm, (b) and (c) both 1 nm. Reprinted with permission from [152]. Copyright (2022) American Chemical Society.

single atomic defect) during experimental operation, therefore a rapid means for detecting and classifying atomic coordinates is of critical importance. In figure 5(b), a single layer of graphene is shown, where the atomic coordinates given by a trained deep ensemble neural network [152, 155] are shown. Graph analysis is performed given the coordinates, from which vacancy defect structures like a Stone-Wales [156] (or, 5–7 ring) defect strongly stand out. As a separate example, in figure 5(c), HAADF-STEM image of MoS₂ is shown also with its detected coordinates shown, where completely different defect structures can be identified and clustered based on, for instance, nearest neighbor distances or relative intensities. In the systems where the electron beam can create atomic defects, each image contains information on the position of multiple atoms in the system in the quasi-equilibrium state. The possible configurations are ultimately determined by the force fields between the atoms. Hence, observations of the atomic configurations allow (in principle) probabilistic reconstruction of the force fields between the atoms.

Similar insight can be derived from mesoscale SPM images. For example, it is the formation of domains that makes of ferroelectrics important for applications in sensors, actuators, capacitors, gates, and information technology devices [145, 158–161]. Therefore, exploring the mechanisms of the ferroelectric domain and domain wall dynamics has been a long-standing problem [162–165]. A ferroelectric domain is an area of aligned spontaneous electric polarization, usually an area in nanoscale or microscale levels. Piezoresponse force microscopy has emerged as a powerful tool for the visualization of ferroelectric domains



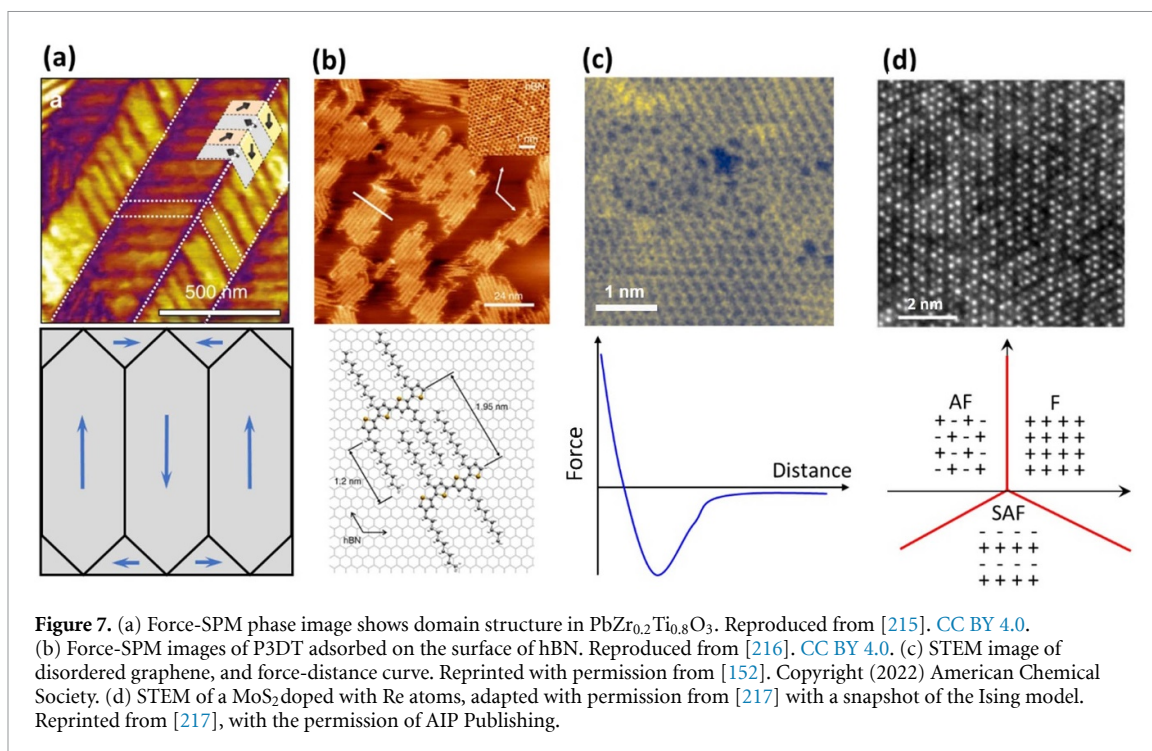
and explore associated physical mechanisms. Domain walls—boundaries between domains with different polarization orientations—with the discontinuity of polarization often present complicated properties due to the presence of defects, disorder, and complex histories. In addition, owing to the polar nature of ferroelectric materials, the bound charge may result in charged walls, leading to strong coupling between domain wall and semiconducting properties [166, 167]. Therefore, domain wall behavior and switching mechanisms at the nanoscale level, including wall motion, domain nucleation, and pinning, root a broad spectrum of properties and applications of ferroelectric materials.

Shown in figure 6 is an example of a ferroelectric domain and its dynamics imaged by piezoresponse force microscopy. The raw phase and amplitude are extracted from a video that indicates continuous domain evolution under applied bias. Here the domain walls can be identified by a Canny filter or DCNN model, as shown in figures 6(c) and (d), which indicated segmented ferroelectric walls and ferroelastic walls, respectively. Then, the identified domain wall videos allow the exploration of relevant mechanisms by ML analysis. Indeed, in our earlier study, using rotationally invariant autoencoders we discovered factors affecting the ferroelectric wall dynamics [168].

2.7. Theory-in-the-loop

For atomic systems, one can run simulations and obtain result while the experiment is still running. For example, as shown in figure 7, we can observe existing structure or assemble a structure in the STEM environment on the atomic level, but we cannot really measure its electronic/magnetic properties (since it requires probing states at the Fermi level). So, we can use theory to get an idea of how ‘interesting’ or ‘useful’ these structures are and to help identify the next step (e.g. what shall we passivate our vacancies with to make sure they remain stable, etc).

Using DL surrogate models, one can in principle allow these behaviors in real-time, opening a connection to physical DL models. Physical models constructed using first-principles theory to quantum Monte Carlo and finite-element methods, spanning over quantum-mechanical to continuum scales, carry a multitude of information on structural, thermodynamic, and electronic properties of materials. Hence, introducing knowledge gathered from theoretical simulations can guide DL models to become more physics(science)-informed. Such understandings from simulations also help in decision-making process by underpinning causal hypothesis-driven [169, 170] mechanisms. The elaborative nature of theoretical models assists in establishing structure-property relations for which structured causal networks work well to answer associational interventional and counterfactual questions. Policies to draw inferences on effects of specific



variables on functionalities of interest in a simulated environment can eventually be translated into active discovery frameworks for automated experiments. It is possible to even design active learning strategies that draw prior information [171] based on theoretical models for exploring both existing as well as unknown chemical spaces.

Be it the DL models, or causal networks or physics-informed optimization schemes for active learning or all combined, there are certain challenges to bringing ‘theory-in-the-loop’ applicable to all of frameworks. Although significant progress has been achieved over the years in high-performance computing, stark differences in time-, length scale and latency still exist between running an experiment versus performing one simulation representing the same scenario. For example, density-functional theory (DFT) and molecular dynamics simulations can optimally generate system sizes of Å to nanometer scales. The length of time required to perform one such simulation that is up to a few microseconds long can be multiple CPU hours, whereas STEM images are typically available at a fraction of a second. Semi-empirical methods such as density functional based tight binding and the extended tight binding method, enable simulations of large systems and relatively long timescales at a reasonable accuracy. These are considerably faster for typical DFT *ab initio* [172], but still takes a few hours to converge. Defining regions of interest in an automated fashion followed by performing simulations on those ‘patches’ is a suitable alternative to draw inferences on a partial or entire system [173]. These can be achieved by multiple means, such as finding optimized fit-to-experiment structures via forward modeling and scalable data analyses.

As a result of these considerations, the typical question of ‘what to simulate’ depends heavily on the available constraints—computational resources, the accuracy required for the particular task (which can be very minimal for explorative work, or may need to be very accurate to match closely with experimental observables), and the specific functionalities explored. For instance, a person investigating ferroelectric domain dynamics would likely want to simulate the domain evolution under applied fields via phase-field models, and in real-time, adapt the phase-field model parameters to both the boundary conditions (surface polarization structure observed, and electrical and elastic boundary conditions of the sample) to produce large set of simulations that can guide the operator in what possible meta-stable states could be observed, and potentially how to reach them. We note that this is possible in mesoscale systems where the partial differential equations (PDE) is given, and relatively well-understood, but in e.g. atomic cases the situation is more complicated and may need a mix of models across length/time scales.

Overall, frameworks bringing microscopic experiment and theory together in real time require capability for rapid exploration of features, analyses of images, quantification of associated uncertainties and most importantly, simulations performed in reasonable (comparable to experiment run time) timescale to establish an active feedback loop for guiding the next set of experiments.

2.8. Rewards

A key but often overlooked aspect of an automated experiment is the reward. Interestingly, on the domain side, it is often understood implicitly and incorporated both in the type of research and experiment planning. Basic science is associated with exploration, e.g. the discovery of new behaviors or finding generalized descriptions of the observed behaviors. Applied science targets exploitation, e.g. optimizing a specific aspect of materials functionality. However, even basic science has the exploration limited to a specific region of knowledge space, or constrained to a strong hypothesis, to make the scientific process tractable. The reward in this case is refining or falsifying the hypothesis. The combination of these elements gives rise to the current paradigm of hypothesis-driven science, allowing for serendipitous discovery.

Comparatively, the reward in ML is usually much more well-defined. For example, in computer games, it is a win. In automated driving it is to get from location A to location B in the least amount of time, with minimal use of fuel, and avoiding accidents and road rule-breaking. Hence, while multiple types of rewards, including external prescribed rewards, curiosity, or empowerment in reinforcement learning, and strategies to balance multiple rewards in multi-objective optimization, have been developed, the reward structure is usually straightforward and is derived externally or from simple information-theoretical criteria.

Microscopy experiments are associated with a much larger range of possible rewards, defined locally within the experimental context or more broadly with respect to the chosen project or experimental workflow. For example, for instrumentation and technique development, the reward is often the increase of spatial and energy resolution [174] or minimization of the beam damage. Of course, time is also a reward, in the sense that increased throughput of imaging is also desirable from an efficiency standpoint. In fact, for the microscopy community, this is often the primary factor driving the development of the microscope over the technological cycle, or tuning the existing system during the experiment. At the same time, this reward is of less interest for domain communities, that aim to explore specific materials or quantify specific aspects of materials behaviors.

Implicitly, the reward structure drives the experiment led by a human operator, and the reward structure and time budget control the experiment planning. Correspondingly, microscopy offers an extensive playground for planning multi-level workflows, with resolution and stability optimization as a lower-level goal and the discovery and experiment planning as a higher-level goal. The former task is defined within a well-defined parameter space, for example the tuning of the scanning probe or electron microscope. At the same time, domain specific tasks offer an interesting scientific and sociological challenge. Can these models and workflows be made machine interpretable? Similarly, given these domain specific rewards, can ML algorithms discover new strategies for tuning the microscope to achieve them?

3. Current state of ML in microscopy

The first applications of modern DL techniques to scanning probe [175] and scanning transmission electron [53] microscopy were demonstrated in 2017 for the off-line data analysis. Since then, there has been a growing interest in using deep and ML to automate routine analysis of microscopy data (such as atom/defect/particle finding) [8, 155, 176–179], learning symmetries [180], and to extract physically meaningful latent parameters explaining high-dimensional observations [138, 143, 144, 157, 181–183]. More recently, the pre-trained DL models were used to identify on-the-fly objects of interests, such as domain walls and atomic defects, in automated experiments in scanning probe and electron microscopy.

However, a significant and well-known limitation of DL models pre-trained on static datasets is that they often fail to generalize OOD, demonstrating poor performance (or even outright failure) when applied to data from outside the domain of training examples. The examples of OOD data in microscopy are images acquired under different acquisition parameters, presence of contaminants on the surface, sample damage due to e-beam irradiation, or change in the probe state due to unexpected interaction with a sample surface. While a combination of ensemble learning and iterative re-training partially addresses this issue for atom-resolved data [155], it does not solve the problem entirely. Furthermore, the approach based on pre-trained DL models assumes that we always know what structures are ‘interesting’, while in many cases this is exactly what we want to discover.

One way to overcome the above limitations is to replace model(s) pre-trained on a static dataset with a model actively interacting with a data generation process. Recently, our group have used an *active* learning approach for learning a probabilistic structure-property relationship on-the-fly and using this ‘knowledge’ to sample next measurement locations [184, 185]. This approach is based on the idea that we often know what physical behavior (functionality) we are interested in and that it is encoded in spectra, and we want to identify (local) structural features where this behavior is maximized. We also want to achieve this with as few measurements as possible. The structure-property relationship of interest is then approximated at each step

via a deep kernel learning model [186], which is used to sample the next measurement locations expected to host the physical behavior of interest.

4. Way forward

The considerations in sections 2 and 3 suggest that there is a tremendous opportunity for research and development at the interface between ML and microscopy. While the potential benefits for the microscopy community *per se* are by now well realized and explored in multiple publications and opinion pieces [187–189], the opportunities for ML community in microscopy domain are much less realized. Below, we summarize the classical components for the developments of an ML ecosystem, but aim to highlight the potential benefits and outcomes from ML side.

4.1. Data repositories: useful but limited

The initial success of ML research was based on (and is continuously measured against) the static datasets, the proverbial big data. While the definitions of big data evolved with time, it is now generally understood that it will refer to the data sets that is sufficient to sample the distribution from which the data is drawn, for example allow training of a generative model. Similar to many other fields, initial effort in microscopy was also dedicated towards the development of supervised models [53, 175]. By analogy with other fields, it has been proposed that large-scale data repositories can be universally useful. Some attempts to create these in microscopy have been made [190, 191].

We argue that this may not be the case, at least not in the direct form. While training examples are definitely useful, they cannot represent all the ways the microscope can work. Even in STEM, the microscope variability can be very significant between different experiments, and the phenomena such as beam distortion due to the incomplete correction can be mixed with (as an example), the atomic column shape distortion due to octahedra tilts in the beam direction [192, 193]. For the simplest example, all ideal crystalline materials with the same lattice will look the same under a microscope (down to lattice parameter). In SPM, the variability of the possible objects is even larger and magnetic domains, ferroelectric domains, polar regions in nanomaterials, or grains in polycrystalline materials can give rise to very similar contrasts.

While at the first glance counterintuitive, we argue that this is simple because the imaging methods are convolution of material and detection system. For the latter, detection system can have variabilities and modalities that far exceed that of the material variability (all crystalline materials look the same under the microscope). Comparatively, all vision tasks are predominantly based on human eye, or compared to the human eye baseline (e.g. in medical imaging analysis is verified via human labeled data sets/experts). Correspondingly, the value of the bespoke image databases can be minimal unless images are augmented by the detailed metadata on microscope parameters and sample composition and history.

At the same time, microscopy creates unique opportunities for the development of novel ML methods. For example, the image plus microscope metadata pairs offers ideal model systems for the invariant risk minimization problems due to the presence of strong inferential biases. In other words, unlike many other computer vision problems, it is well known that the answer exists and is unique. This further opens the pathway for developing ML methods that allows for problem-specific augmentation (variability of the imaging system), encapsulate and aim to discover data generation process, etc.

4.2. Hyper language: tuning and manipulations

On a fundamental level, every microscopy experiment consists of a carefully orchestrated sequence of tasks, often with inner and outer optimization loops, and decision making at particular steps. This appears obvious when stated but is in fact so ingrained in the standard microscope operator's daily routine that it is rarely formulated in explicit form. However, in order to enable fully autonomous microscopes, we need to develop a language that can represent every low-level task that a microscope system can perform. This can be thought of as the microscopy equivalent of the recently developed 'XDL' programming language for chemical synthesis, which is a Turing-complete language that is in principle able to represent all types of chemical reactions [194].

Naturally, the low-level microscopy functions would be operations such as moving the stage and/or tip to a certain location, scanning along a trajectory, enabling and disabling feedback systems, measuring along different channels, changing local and global excitation parameters, and so forth, and there would additionally be 'decision blocks' for where decisions would be made about the next course of action (either by a human or algorithm).

As an example, consider a simply SPM experiment of measuring the roughness of a sample. The basic steps would be to initiate the tapping mode operation (including the tuning steps), scan regions of a particular size that are free from large debris and move the scan window after each scan to collect more

images and thus statistics. Then the tip should be withdrawn, and the images analyzed with a standard topographic analysis to determine the average sample roughness. This could easily be represented as a sequence of function calls in a script, but it could also be represented by an even more compact representation akin to that of molecules—such as the SELFIES representation introduced by Aspuru-Guzik group [195, 196]. Then, performing an experiment and generating new experiments becomes an exercise in generating new such ‘microscope SELFIES’.

4.3. Beyond human workflows

An interesting and very weakly explored concept in the context of automated experiment is the development of beyond-human workflows. Indeed, currently most of the effort in automated experiment relies on classical human based workflows, implemented via robotic agents. An alternative is combined human and robotics operators orchestrated by central agent.

At the same time, recent advances in reinforcement learning have demonstrated that for closed systems with known rules but complex emergent strategies (e.g. Alpha Go), AI is capable of formulating the strategies beyond human. Interestingly, this leads to the inverse developments when exposure to the AI agents improves human learning. Until now, these developments have been limited to fully closed system, albeit recent progress with large models such as DALLÉ make in the art field suggest that the situation is likely to change.

Similar developments, albeit with a certain time delay, are now proceeding in the field of automated labs. To date, the vast majority of automated or hybrid labs utilize the human made workflows derived from human experts or via literature mining [197–201]. In many cases, ML is used to optimize specific steps such as yield of specific product via optimization in low-dimensional continuous parameter space and implemented via continuous flow reactors, microfluidic systems, or pipetting robots [202]. However, proponents such as Symes *et al* [203] now explore pathways for full chemistry design with 3D printing of components. That said, exploration of the chemical space of e.g. small molecules is an exceptionally complex problem due to its immense size and non-differentiability [204]. For the experimental world, this often limits discovery space to the molecules that can be synthesized from a limited set of precursors and reactions.

Comparatively, microscopy offers a large, but still much smaller realm for exploration. Systems such as atomic manipulation by electron beam or STM probe or ferroelectric domain manipulation by Piezoresponse Force Microscopy (PFM) offer the richness of unknown physical behaviors, a multitude of possible metastable states, and a variety of ways to reach them. However, these systems come with a finite range of possible actions and actions that are fully available within the specific system. As such, this can be an ideal toy system for exploring ML creation of workflows starting from optimization of human based ones but evolving to discover new ones, via empowerment learning or other techniques. Interestingly, even deep kernel learning discovery of the plasmonic EELS signals already presents beyond human workflow.

4.4. Learning physics

Generally, a physics-based solution, in circumstance where the physics is computable and understood, is much preferable to a purely data-driven approach, due to efficiency of learning. In most cases, the physics-models may have zero, one or a few parameters to be estimated. As an example, consider that the standard ‘cartpole’ reinforcement learning environment, where the aim is to move the cart left or right to balance a pole in the center of the cart without it tipping over, is solvable using deep reinforcement learning techniques, which can take 100s of episodes, or, it can be solved in 5 lines of code without any RL, given knowledge of Newtonian mechanics. Of course, in a real scientific problem, we may have to infer the underlying model from noisy measurements of the system, often sparsely in space and/or time.

Notably, the ML community is developing methods that can attempt to recover the underlying equations that generate the observations measured, through techniques such as symbolic regression. These are exemplified by the approach of the ‘AI Feynman’ proposed by Wu and Tegmark [205–207], amongst others [208, 209]. In general, the applications of these techniques to microscopy are still very limited. For instance, it would be instructive to learn the force-fields from an atomically resolved STEM video of an amorphous material being crystallized by the electron beam, or of directly determining the energy of barriers for dopants imaged moving through a 3D lattice. Although much simpler in the static case [210], the extension to a fully dynamic system naturally tends to lead towards the need for large volumes of data, and careful tuning of thermodynamic variables to efficiently determine the likelihood of proposed models describing the dynamics.

4.5. Causal processes, interventions, and counterfactuals

Traditional causal discovery approaches assume that the observations are structured into elementary units representing random variables connected by a causal graph. In some cases, we may treat different modalities in multimodal SPM as such units. Practically, an image in each observational channel is split into patches at

selected coordinates, followed by an established post-processing analysis to form appropriate structural and functional descriptors associated with such units. Then, one can deploy causal discovery models for finding causal links between various structural and functional parameters.

At the same time, in many experiments, the observations cannot be directly mapped onto a causal graph, and one must first uncover ('learn') in an unsupervised manner low-dimensional, high-level causal variables from high-dimensional, low-level observational data to allow for causal discovery. The emerging field of causal representation learning [211] aims at addressing this problem through an extension of the deep latent variables (correlative) models. The goal of deep latent variable models, such as (variational) autoencoder ((V)AE), is to explain high-dimensional observations by a small number of unobserved (latent) variables [212]. Recently, VAEs were used to uncover physical order parameters in lattice models and dynamic atomically resolved imaging data [182]. However, the latent variables in those VAE setups did not support intervention and reasoning. Recently, an approach based on self-supervised learning with data augmentation has been proposed to overcome this limitation of traditional deep latent variable models. In this approach, data augmentation acts as a counterfactual under a hypothetical intervention, allowing to separate content from style. We argue that in microscopy, one could replace data augmentations by changing various imaging parameters during the experiment to disentangle the real physical variables and use them for further causal discovery.

5. Future

Finally, we aim to summarize what these developments can bring to the classical physics, chemistry, and materials science fields, including both the extant ones, as well as new areas that may emerge based on these developments. While the list below is undoubtedly partial and new applications can be expected to be ideated and implemented, it nonetheless suggests an extremely broad potential for emerging science at the interface between ML and microscopy.

5.1. Learn physics

Traditional physics describes the structure of crystalline materials as ordering of building blocks (elementary unit cells) in a periodic lattice. These descriptors are conveniently matched to the information that can be obtained from scattering methods, and in turn allows exploring materials structure in the Fourier domain. However, many important materials such as highly defective solids, structural, spin, and cluster glasses, and ferroelectric relaxors do not have periodic structures [213, 214]. While their structures can be visualized via electron microscopy, it is not clear what information can be derived from the knowledge of the atomic positions. Similar opportunities arise in the context of the electron microscopy of the extended defects and multiphase materials, STM of materials with complex topological and electronic orders, and even conventional SPM of ferroelectric and ferromagnetic domains and materials microstructures. The microscopy images are often spectacular, contain information on the specific aspects of materials structures and functionality, but extracting this information is non-trivial.

One approach to use these data is to train the generative statistical models such as VAEs and GANs, that can provide multiple examples of the microstructure from a smaller number of observations. However, such models typically extrapolate very poorly, meaning that the data acquired for specific set of conditions will not generalize to a different one. At the same time, the structural complexity of many materials strongly depends on temperature. For example, in materials undergoing order-disorder phase transition, correlation length (and hence the size of structural descriptor) diverges in the vicinity of phase transitions.

Comparatively, the complexity of observed structures and functionalities of solids emerges from relatively simple interactions. For atomic distributions in solid solutions and certain magnetic systems, these can be the variants of an Ising Hamiltonian, representing the interactions in the system via pairwise exchange integrals. For structurally amorphous systems, the system can be described via force fields acting between the atoms. Note that the modeling of the systems behavior from known force fields or Hamiltonians is the foundational part of computational physics. However, relevant materials parameters are usually derived ad hoc, or fit from a large body of macroscopic observations.

Hence, the obvious challenge for the ML in physics is probabilistic learning of the atomic interactions on the level of lattice models, force fields from the observational data, or phase field free energies. These physical generative models can provide insight into the fundamental physics and chemistry of materials.

5.2. New chemistry

Synthetic organic chemistry is a cornerstone of fields ranging from drug and catalyst design to polymers. Applications of ML in chemical synthesis to understand the relationships between molecule structure and biological functionality, predict the synthesizability, and predict possible synthetic pathways from available

components is one of the most rapidly developing scientific directions [218, 219], via graph networks [5, 220], etc. The rapid emergence of cloud laboratories provides a real-world implementation for *in-silico* developments. However, by its very nature chemical space is extremely high-dimensional, and exploring even specific classes of compounds often requires access to large number of initial compounds and ligands.

The emergence of imaging methods allows us to merge chemistry and characterization. The first definitive example of this synergy is the overlap between CryoEM and biology, where local structural information on biomolecules is now being used to design biological targets, etc. Electron microscopy methods can provide similar insights into the chemistry of solids. For example, STEM data can be used derive a library of possible defects in solids [176]. We can also extend this to explore whether we can chemically modify them and build reaction schemes on top of it as a pathway towards better catalysts, etc.

5.3. Nanophotonics and quantum optics

The wavelength of light, about half a micron, fundamentally limits its use in non-biological nanotechnology—hence the primary adoption of electrical signals to this day. Currently, using extreme ultraviolet exposure methods, the semiconductor industry has pushed the limits of creating nanoelectronic devices down to the 3~nm node length, but such devices all rely on the use of electrons or holes for computations and information transfer, which inefficiently generate heat and is relatively slow. On the other hand, light propagates at the ultimate speed (of light!) and because photons do not have a rest mass, it does not lose energy to Joule heating. However, it cannot be easily confined to volumes smaller than the diffraction limit, or about 250 nm.

Plasmonics offers a tempting route to bridge the macro and nano scales with light by exploiting a quasiparticle known as the plasmon, which is light coupled to a collective electron wave at the interface of (usually) a metal and insulator, either propagating across an interface or localized to a small volume. The former is known as a surface plasmon polariton (SPP), while the latter is known as a localized surface plasmon resonance (LSPR). Intriguingly, a strong enhancement of the light can occur when it is coupled to the surface plasmon [221], making for a useful attribute. An idea for information transfer is that an incoming light source on the order of millimeters in size is focused down by ordinary means close to the diffraction limit and then partially transformed into a propagating wave (SPP) that is tens of nanometers in size, therefore the nature of the plasmon acts as a bridge between these two different length scales. Other applications take advantage of the LSPR in order to ensure extremely localized (and enhanced) light or heat sources. For example, in some cancer treatments, small nanoparticles may be placed in the body and then illuminated by a light source (to which the human body is transparent), which in turn excite LSPRs in the nanoparticles that act as tiny, localized heat sources to attack cancerous cells [222].

In terms of quantum optics, controlling the location, energy, and spatial extent of an emitter is of utmost importance, particularly in the rapidly developing field of quantum computing. Plasmonic nanoparticles can act as localized emitters that can in turn excite, for instance, a dye molecule which itself may be a single photon emitter. Alternatively, and perhaps of more utility, point defects in 2D materials like boron nitride [223] can also act as quantum emitters, and the electron beam can potentially be used to both create and probe such single defects.

The STEM then offers a unique and well-suited platform for nanophotonics [224]. From the point of view of measuring the light-matter interaction's energy dependence in space at the nanometer scale, STEM-EELS is the only currently available tool to spatially resolve the energy dependence of plasmon modes. A perfect corollary to the nano and atomic scale measurement capability is the ability of the STEM to geometrically sculpt or chemically modify a two- or three-dimensional solid such that it suits a specific plasmonic need, which can then be measured in STEM-EELS almost immediately.

However, realizing the full potential of STEM for nanophotonics and quantum optics requires solving the inverse design problem, i.e. what geometries and positioning/type of defects are necessary to induce the required functional responses. This space is ripe for exploration with ML methodologies [152], and one can envision both generative as well as reinforcement learning approaches to tackle this inverse problem, ideally incorporating simulations where necessary to reduce data requirements.

5.4. Atomic fabrication

Fabricating matter atom by atom has long remained the dream of scientific community. In his final note, Richard Feynman had famously written 'what I cannot create, I do not understand' [225]. It is clear that atomic fabrication [226–229] at scale is a pathway for quantum devices, new electronics, and serendipitous discoveries in multiple areas of science and technology.

Currently, both STM and STEM can be used to manipulate matter on the atomic scale, as have been demonstrated by multiple groups over time [67, 230, 231]. However, while the enabling instrumentation and the engineering controls are currently available, the operation of the microscopes is preponderantly enabled

by human operator. This is a clear area where ML can make a tangible difference. For STM, the relevant tasks include tip conditioning and optimization, rapid sampling of surface electronic structure, investigating the electronic properties of individual atomic groups, physics discovery based on desired electronic signatures, and learning the rules of atomic manipulation towards specific functionalities and structures. Similar problems emerge in the STEM context, with added complexity of the higher stochasticity of the system.

6. Our vision

The last decade has seen the exponential growth of ML methods that by now have become an inseparable part of multiple scientific and applied domains. However, the emerging challenge in ML is a transition away from static data sets towards exploring active data generation processes and enabling autonomous systems, from automated cars to chemical design and synthesis. Each of these fields encounter problems related to the inability of classical ML to generalize and extrapolate. At the same time, many of these fields own specific physical interferences, are many hidden degrees of freedom, as well as exogenous variables.

We argue that microscopy platforms that operate in very simplified (compared to the macroscopic world) physical environments, and that are associated with strong physical biases and low risks, are ideal toy systems for deployment of active learning algorithms, learning physics based digital twins, and many other applications. While these tools have traditionally been perceived to be highly specialized, over the last two decades microscopes have evolved to become significantly more user-friendly, with effective APIs and easy to use systems capable of visualizing matter from atomic to micron levels. Further, they are capable of a broad spectrum of actions, ranging from manipulation of nanoparticles and nanowires, to inducing local chemical reactions, and moving single atoms and making and breaking chemical bonds. From the domain perspective, the missing component is the algorithms that enabled controlled manipulation towards specific materials objectives. At the same time, from the ML perspective, these domains offer a challenging but bounded environment with characteristics ideal for development and deployment of active learning and reinforcement learning methods. In contrast to many other traditional environments (e.g. for robotics), these systems are associated with a reduced set of commands, low risk, and well understood physical biases and causal relationships, often with minimal exogenous effects. At some point, a Chat-GPT like system can be deployed at the edge to assist with operating a physical instrument such as a microscope. For example, it can be trained on a large dataset of instructions and common queries related to the operation of the instrument and measurement of different samples. This can allow the model to generate natural language responses to user queries and provide step-by-step instructions for operating the instrument for different type of samples/experiments. In addition, a ChatGPT-like system trained on scientific literature and relevant materials databases can dramatically accelerate the traditional hypothesis-experiment cycle by using AI to automatically generate hypotheses and integrate them into active learning models that ‘drive’ characterization experiments in microscopy. In other words, microscopy is all you need (to start).

Data availability statement

No new data were created or analyzed in this study.

Acknowledgments

This research is sponsored by the INTERSECT Initiative as part of the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy under Contract DE-AC05-00OR22725. This effort was primarily supported by the center for 3D Ferroelectric Microelectronics (3DFeM), an Energy Frontier Research Center funded by the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences under Award Number DE-SC0021118. This research was supported by the Center for Nanophase Materials Sciences (CNMS), which is a US Department of Energy, Office of Science User Facility at Oak Ridge National Laboratory. The authors are grateful to Professor J Agar (Drexel) and K Brown (Boston University) and Marc Bellemare (Google Brain) for early comments on the manuscript.

ORCID iDs

Rama Vasudevan  <https://orcid.org/0000-0003-4692-8579>

Yongtao Liu  <https://orcid.org/0000-0003-0152-1783>

Ayana Ghosh  <https://orcid.org/0000-0002-0432-3689>

Maxim Ziatdinov  <https://orcid.org/0000-0003-2570-4592>

References

- [1] Heravi E J, Aghdam H H and Puig D *CCIA* pp 163–8
- [2] Deng J et al 2009 ImageNet: a large-scale hierarchical image database ‘CVPR09’ *IEEE Computer Society (Miami, Florida, USA)* pp 248–55
- [3] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [4] Schmidhuber J 2015 Deep learning in neural networks: an overview *Neural Netw.* **61** 85–117
- [5] Szymkuc S, Gajewska E P, Klucznik T, Molga K, Dittwald P, Startek M, Bajczyk M and Grzybowski B A 2016 Computer-assisted synthetic planning: the end of the beginning *Angew. Chem., Int. Ed.* **55** 5904–37
- [6] Higgins K, Valletti S M, Ziatdinov M, Kalinin S V and Ahmadi M 2020 Chemical robotics enabled exploration of stability in multicomponent lead halide perovskites via machine learning *ACS Energy Lett.* **5** 3426–36
- [7] Higgins K, Ziatdinov M, Kalinin S V and Ahmadi M 2021 High-throughput study of antisolvents on the stability of multicomponent metal halide perovskites through robotics-based synthesis and machine learning approaches (arXiv:2106.03312)
- [8] Gordon O M, Hodgkinson J E A, Farley S M, Hunsicker E L and Moriarty P J 2020 Automated searching and identification of self-organized nanostructures *Nano Lett.* **20** 7688–93
- [9] Krizhevsky A, Sutskever I and Hinton G E 2012 ImageNet classification with deep convolutional neural networks *Advances in Neural Information Processing Systems (NIPS 2012)* vol 25
- [10] Kearnes S, McCloskey K, Berndl M, Pande V and Riley P 2016 Molecular graph convolutions: moving beyond fingerprints *J. Comput. Aided Mol. Des.* **30** 595–608
- [11] Vaswani A et al 2017 Attention is all you need *Advances in neural information processing systems* vol 30
- [12] Kingma D P and Welling M 2013 Auto-encoding variational bayes (arXiv:1312.6114)
- [13] Kingma D P and Welling M 2019 An introduction to variational autoencoders *Found. Trends Mach. Learn.* **12** 307–92
- [14] Cerri O, Nguyen T Q, Pierini M, Spiropulu M and Vlimant J R 2019 Variational autoencoders for new physics mining at the large hadron collider *J. High Energy Phys.* **JHEP05(2019)036** 36
- [15] Goodfellow I et al 2014 Generative adversarial nets vol 27
- [16] Bobkov Y V, Walker W B III and Cattaneo A M 2021 Altered functional properties of the codling moth Orco mutagenized in the intracellular loop-3 *Sci. Rep.* **11** 1–16
- [17] Grant-Jacob J A, Praeger M, Loxham M, Eason R W and Mills B 2020 Lensless imaging of pollen grains at three-wavelengths using deep learning *Environ. Res. Commun.* **2** 075005
- [18] Nguyen T, Xue Y, Li Y, Tian L and Nehmetallah G 2018 Deep learning approach for Fourier ptychography microscopy *Opt. Express* **26** 26470–84
- [19] Qiao C, Li D, Guo Y, Liu C, Jiang T, Dai Q and Li D 2021 Evaluation and development of deep neural networks for image super-resolution in optical microscopy *Nat. Methods* **18** 194–202
- [20] Schrittwieser J et al 2020 Mastering Atari, Go, chess and shogi by planning with a learned model *Nature* **588** 604
- [21] Silver D et al 2018 A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play *Science* **362** 1140
- [22] Berner C et al 2019 Dota 2 with large scale deep reinforcement learning (arXiv:1912.06680)
- [23] Andrychowicz O M et al 2020 Learning dexterous in-hand manipulation *Int. J. Rob. Res.* **39** 3–20
- [24] Bellemare M G, Candido S, Castro P S, Gong J, Machado M C, Moitra S, Ponda S S and Wang Z 2020 Autonomous navigation of stratospheric balloons using reinforcement learning *Nature* **588** 77–82
- [25] McCulloch W S and Pitts W 1943 A logical calculus of the ideas immanent in nervous activity *Bull. Math. Biophys.* **5** 115–33
- [26] Widrow B and Lehr M A 1990 30 Years of adaptive neural networks: perceptron, Madaline, and backpropagation *Proc. IEEE* **78** 1415–42
- [27] Rosenblatt F 1961 *Principles of Neurodynamics. Perceptrons and the Theory of Brain Mechanisms* (Buffalo, NY: Cornell Aeronautical Lab Inc)
- [28] Santos M, Ludermir T, Santos F, de Melo C P and Gomes J 1970 Artificial nose and data analysis using multi layer perceptron *WIT Trans. Inf. Commun. Technol.* **22**
- [29] Rumelhart D E, Hinton G E and Williams R J 1986 Learning representations by back-propagating errors *Nature* **323** 533–6
- [30] Widrow B, Gupta N K and Maitra S 1973 Punish/reward: learning with a critic in adaptive threshold systems *IEEE Trans. Syst. Man Cybern.* **3** 455–65
- [31] Sutton R S and Barto A G 2018 *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press)
- [32] Watkins C J C H 1989 Learning from delayed rewards
- [33] Hochreiter S and Schmidhuber J 1997 Long short-term memory *Neural Comput.* **9** 1735–80
- [34] Dally W J, Keckler S W and Kirk D B 2021 Evolution of the graphics processing unit (GPU) *IEEE Micro* **41** 42–51
- [35] Deng L 2012
- [36] Krizhevsky A and Hinton G 2009 Learning multiple layers of features from tiny images
- [37] Schatzki L, Arrasmith A, Coles P J and Cerezo M 2021 Entangled datasets for quantum machine learning (arXiv:2109.03400)
- [38] Prabhat M et al 2017 AGU fall meeting abstracts IN13E-01
- [39] Hollis D, McCarthy M, Kendon M, Legg T and Simpson I 2019 HadUK-Grid—a new UK dataset of gridded climate observations *Geosci. Data J.* **6** 151–9
- [40] Racah E et al 2017 Extremeweather: a large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events *Advances in Neural Information Processing Systems* vol 30
- [41] Foppiano L et al 2021 SuperMat: construction of a linked annotated dataset from superconductors-related publications *Sci. Technol. Adv. Mater.* **1** 34–44
- [42] Borg C K H, Frey C, Moh J, Pollock T M, Gorse S, Miracle D B, Senkov O N, Meredig B and Saal J E 2020 Expanded dataset of mechanical properties and observed phases of multi-principal element alloys *Sci. Data* **7** 1–6
- [43] Choudhary K et al 2022 Recent advances and applications of deep learning methods in materials science *npj Comput. Mater.* **8** 1–26
- [44] Marcus G and Davis E 2019 *Rebooting AI: Building Artificial Intelligence We Can Trust* (Vintage)
- [45] Brown T et al 2020 Language models are few-shot learners *Adv. Neural Inf. Process. Syst.* **33** 1877–901
- [46] Quach K 2020
- [47] Amodei D et al 2016 Concrete problems in AI safety (arXiv:1606.06565)

- [48] Nguyen A, Yosinski J and Clune J 2015 Deep neural networks are easily fooled: high confidence predictions for unrecognizable images *Proc. IEEE conference on Conf. on Computer Vision and Pattern Recognition* pp 427–36
- [49] Gulrajani I and Lopez-Paz D 2020 In search of lost domain generalization (arXiv:2007.01434)
- [50] Zech J R, Badgeley M A, Liu M, Costa A B, Titano J J and Oermann E K 2018 Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study *PLoS Med.* **15** e1002683
- [51] Nagarajan V, Andreassen A and Neyshabur B 2020 Understanding the failure modes of out-of-distribution generalization (arXiv:2010.15775)
- [52] Beery S, Van Horn G and Perona P 2018 Recognition in terra incognita *Proc. European Conf. on Computer Vision (ECCV)* pp 456–73
- [53] Ziatdinov M, Dyck O, Maksov A, Li X, Sang X, Xiao K, Unocic R R, Vasudevan R, Jesse S and Kalinin S V 2017 Deep learning of atomically resolved scanning transmission electron microscopy images: chemical identification and tracking local transformations *ACS Nano* **11** 12742–52
- [54] Chopra S, Hadsell R and LeCun Y 2005 Learning a similarity metric discriminatively, with application to face verification 2005 *IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR'05)* vol 531 pp 539–46
- [55] Zbontar J, Jing L, Misra I, LeCun Y and Deny S 2021 Barlow twins: self-supervised learning via redundancy reduction *Int. Conf. on Machine Learning* pp 12310–20
- [56] Lakshminarayanan B, Pritzel A and Blundell C 2017 *Proc. 31st Int. Conf. on Neural Information Processing Syst.* (Long Beach, CA: Curran Associates Inc.) pp 6405–16
- [57] Arjovsky M, Bottou L, Gulrajani I and Lopez-Paz D 2019 Invariant risk minimization (arXiv:1907.02893)
- [58] LeCun Y 2022 A path towards autonomous machine intelligence version 0.9. 2 (Accessed 27 June 2022) (<https://doi.org/10.1016/j.neunet.2022.03.037>)
- [59] Vasudevan R K, Ziatdinov M, Vlcek L and Kalinin S V 2021 Off-the-shelf deep learning is not enough, and requires parsimony, Bayesianity, and causality *npj Comput. Mater.* **7** 6
- [60] Fujinuma N, DeCost B, Hatrick-Simpers J and Lofland S E 2022 Why big data and compute are not necessarily the path to big materials science *Commun. Mater.* **3** 1–9
- [61] Liu Y, Fields S S, Mimura T, Kelley K P, Trolier-mckinstry S, Ihlefeld J F and Kalinin S V 2022 Exploring leakage in dielectric films via automated experiments in scanning probe microscopy *Appl. Phys. Lett.* **120** 182903
- [62] Pennycook S J and Nellist P D 2011 (New York: Springer)
- [63] Yankovich A B, Berkels B, Dahmen W, Binev P, Sanchez S I, Bradley S A, Li A, Szlufarska I and Voyles P M 2014 Picometre-precision analysis of scanning transmission electron microscopy images of platinum nanocatalysts *Nat. Commun.* **5** 4155
- [64] Roccapriore K M, Cho S-H, Lupini A R, Milliron D J and Kalinin S V 2022 Sculpting the plasmonic responses of nanoparticles by directed electron beam irradiation *Small* **18** 2105099
- [65] Varela M et al 2004 Spectroscopic imaging of single atoms within a bulk solid *Phys. Rev. Lett.* **92** 095502
- [66] Kalinin S V, Borisevich A and Jesse S 2016 Fire up the atom forge *Nature* **539** 485–7
- [67] Dyck O, Ziatdinov M, Lingerfelt D B, Unocic R R, Hudak B M, Lupini A R, Jesse S and Kalinin S V 2019 Atom-by-atom fabrication with electron beams *Nat. Rev. Mater.* **4** 497–507
- [68] Dyck O, Kim S, Kalinin S V and Jesse S 2017 Placing single atoms in graphene with a scanning transmission electron microscope *Appl. Phys. Lett.* **111** 113104
- [69] Susi T, Kepaptsoglou D, Lin Y-C, Ramasse Q M, Meyer J C, Suenaga K and Kotakoski J 2017 Towards atomically precise manipulation of 2D nanostructures in the electron microscope *2D Mater.* **4** 042004
- [70] Dyck O, Kim S, Jimenez-Izal E, Alexandrova A N, Kalinin S V and Jesse S 2018 Building structures atom by atom via electron beam manipulation *Small* **14** 1801771
- [71] Jesse S et al 2018 Direct atomic fabrication and dopant positioning in Si using electron beams with active real-time image-based feedback *Nanotechnology* **29** 255303
- [72] Jesse S and Kalinin S V 2011 Band excitation in scanning probe microscopy: sines of change *J. Phys. D: Appl. Phys.* **44** 464006
- [73] Butt H, Cappella B and Kappell M 2005 Force measurements with the atomic force microscope: technique, interpretation and applications *Surf. Sci. Rep.* **59** 1–152
- [74] Binnig G, Quate C F and Gerber C 1986 ATOMIC FORCE MICROSCOPE *Phys. Rev. Lett.* **56** 930–3
- [75] Albrecht T R, Grutter P, Horne D and Rugar D 1991 Frequency-modulation detection using high-q cantilevers for enhanced force microscope sensitivity *J. Appl. Phys.* **69** 668–73
- [76] Grutter P, Liu Y, LeBlanc P and Durig U 1997 Magnetic dissipation force microscopy *Appl. Phys. Lett.* **71** 279–81
- [77] Alvarez T, Kalinin S V and Bonnell D A 2001 Magnetic-field measurements of current-carrying devices by force-sensitive magnetic-force microscopy with potential correction *Appl. Phys. Lett.* **78** 1005–7
- [78] Tanimoto M 1996 Kelvin probe force microscopy for characterization of semiconductor devices and processes *J. Vac. Sci. Technol. B* **14** 1547–51
- [79] Strelcov E, Jesse S, Huang Y-L, Teng Y-C, Kravchenko I I, Chu Y-H and Kalinin S V 2013 Space- and time-resolved mapping of ionic dynamic and electroresistive phenomena in lateral devices *ACS Nano* **7** 6806–15
- [80] Rabe U and Arnold W 1994 Acoustic microscopy by atomic-force microscopy *Appl. Phys. Lett.* **64** 1493–5
- [81] Balke N, Bdikin I, Kalinin S V and Kholkin A L 2009 Electromechanical imaging and spectroscopy of ferroelectric and piezoelectric materials: state of the art and prospects for the future *J. Am. Ceram. Soc.* **92** 1629–47
- [82] Raeder T M et al 2004 High velocity, low-voltage collective in-plane switching in (100) BaTiO₃ thin films *Adv. Sci.* **9** 2201530
- [83] Shin J, Meunier V, Baddorf A P and Kalinin S V 2004 Nonlinear transport imaging by scanning impedance microscopy *Appl. Phys. Lett.* **85** 4240–2
- [84] Garcia R, Martinez R V and Martinez J 2006 Nano-chemistry and scanning probe nanolithographies *Chem. Soc. Rev.* **35** 29–38
- [85] Woo J, Hong S, Setter N, Shin H, Jeon J-U, Pak Y E and No K 2001 Quantitative analysis of the bit size dependence on the pulse width and pulse voltage in ferroelectric memory devices using atomic force microscopy *J. Vac. Sci. Technol. B* **19** 818–24
- [86] Requicha A A G et al 1998 *IEEE Int. Conf. Robotics and Automation (IEEE)* pp 3368–74
- [87] Requicha A A G 2003 Nanorobots, NEMS, and nanoassembly *Proc. IEEE* **91** 1922–33
- [88] Baur C et al 1997 Robotic nanomanipulation with a scanning probe microscope in a networked computing environment *J. Vac. Sci. Technol. B* **15** 1577–80
- [89] Iwata F, Ohashi Y, Ishisaki I, Picco L M and Ushiki T 2013 Development of nanomanipulator using a high-speed atomic force microscope coupled with a haptic device *Ultramicroscopy* **133** 88–94

- [90] Guthold M *et al* 1999 Investigation and modification of molecular structures with the nanoManipulator *J. Mol. Graph.* **17** 187–97
- [91] Falvo M R, Clary G, Helsen A, Paulson S, Taylor R M, Chi V, Brooks F P, Washburn S and Superfine R 1998 Nanomanipulation experiments exploring frictional and mechanical properties of carbon nanotubes *Microsc. Microanal.* **4** 504–12
- [92] Binnig G and Rohrer H 1982 Scanning tunneling microscopy *Helv. Phys. Acta* **55** 726–35
- [93] Binnig G, Rohrer H, Gerber C and Weibel E 1983 7X7 reconstruction on Si(111) resolved in real space *Phys. Rev. Lett.* **50** 120–3
- [94] Binnig G and Rohrer H 1987 Scanning tunneling microscopy—from birth to adolescence *Rev. Mod. Phys.* **59** 615–25
- [95] Eigler D M and Schweizer E K 1990 Positioning single atoms with a scanning tunnelling microscope *Nature* **344** 524–6
- [96] Schofield S R, Curson N J, Simmons M Y, Rueß F J, Hallam T, Oberbeck L and Clark R G 2003 Atomically precise placement of single dopants in Si *Phys. Rev. Lett.* **91** 136104
- [97] He Y, Gorman S K, Keith D, Kranz L, Keizer J G and Simmons M Y 2019 A two-qubit gate between phosphorus donor electrons in silicon *Nature* **571** 371–5
- [98] Kalinin S V *et al* 2016 Big, deep, and smart data in scanning probe microscopy *ACS Nano* **10** 9068–86
- [99] Kalinin S V, Sumpter B G and Archibald R K 2015 Big-deep-smart data in imaging for guiding materials design *Nat. Mater.* **14** 973
- [100] Spurgeon S R *et al* 2021 Towards data-driven next-generation transmission electron microscopy *Nat. Mater.* **20** 274–9
- [101] Holstad T S *et al* 2020 Application of a long short-term memory for deconvoluting conductance contributions at charged ferroelectric domain walls *npj Comput. Mater.* **6** 1–7
- [102] Meyer C, Dellby N, Hachtel J A, Lovejoy T, Mittelberger A and Krivanek O 2019 Nion swift: open source image processing software for instrument control, data acquisition, organization, visualization, and analysis using Python *Microsc. Microanal.* **25** 122–3
- [103] PyJEM (available at: <https://github.com/PyJEM/PyJEM>)
- [104] Pinkard H *et al* 2021 Pycro-manager: open-source software for customized and reproducible microscope control *Nat. Methods* **18** 226–8
- [105] Zhang J X *et al* 2011 Large field-induced strains in a lead-free piezoelectric material *Nat. Nanotechnol.* **6** 97–101
- [106] Ovchinnikov O S, Jesse S and Kalinin S V 2009 Adaptive probe trajectory scanning probe microscopy for multiresolution measurements of interface geometry *Nanotechnology* **20** 255701
- [107] Collins L, Belianinov A, Somnath S, Balke N, Kalinin S V and Jesse S 2016 Full data acquisition in Kelvin probe force microscopy: mapping dynamic electric phenomena in real space *Sci. Rep.* **6** 30557
- [108] Vasudevan R K 2022 2022 *Gordon Research Conf. on Computational Materials Science and Engineering*
- [109] Meyer T R, Ziegler D, Brune C, Chen A, Farnham R, Huynh N, Chang J-M, Bertozzi A L and Ashby P D 2014 Height drift correction in non-raster atomic force microscopy *Ultramicroscopy* **137** 48–54
- [110] Kelley K P, Ziatdinov M, Collins L, Susner M A, Vasudevan R K, Balke N, Kalinin S V and Jesse S 2020 Fast scanning probe microscopy via machine learning: non-rectangular scans with compressed sensing and Gaussian process optimization *Small* **16** 2002878
- [111] Stevens A, Kovarik L, Abellan P, Yuan X, Carin L and Browning N D 2015 Applying compressive sensing to TEM video: a substantial frame rate increase on any camera *Adv. Struct. Chem. Imaging* **1** 1–20
- [112] Ganpule C S, Nagarajan V, Li H, Ogale A S, Steinhauer D E, Aggarwal S, Williams E, Ramesh R and De Wolf P 2000 Role of 90 degrees domains in lead zirconate titanate thin films *Appl. Phys. Lett.* **77** 292–4
- [113] Arnold C 2022 Cloud labs: where robots do the research *Nat. Mater.* **606** 612–3
- [114] Horcas I, Fernández R, Gómez-Rodríguez J M, Colchero J, Gómez-Herrero J and Baro A M 2007 WSXM: a software for scanning probe microscopy and a tool for nanotechnology *Rev. Sci. Instrum.* **78** 013705
- [115] Nečas D and Klapetek P 2012 Gwyddion: an open-source software for SPM data analysis *Open Phys.* **10** 181–8
- [116] Scholder O 2018 scholi/pySPM: pySPM v0.2.16 (Version v0.2.16)
- [117] SciFiReaders
- [118] sidpy 2022
- [119] de la Peña F *et al* 2017 Electron microscopy (big and small) data analysis with the open source software package HyperSpy *Microsc. Microanal.* **23** 214–5
- [120] openNCEM
- [121] LiberTEM
- [122] Collins J T *et al* 2020 Robotic microscopy for everyone: the OpenFlexure microscope *Biomed. Opt. Express* **11** 2447–60
- [123] Mehrabi N, Morstatter F, Saxena N, Lerman K and Galstyan A 2021 A survey on bias and fairness in machine learning *ACM Comput. Surv.* **54** 1–35
- [124] Melcher J, Hu S and Raman A 2008 Invited Article: VEDA: a web-based virtual environment for dynamic atomic force microscopy *Rev. Sci. Instrum.* **79** 061301
- [125] Silver D *et al* 2017 Mastering the game of Go without human knowledge *Nature* **550** 354
- [126] Nicoletti O, de la Peña F, Leary R K, Holland D J, Ducati C and Midgley P A 2013 Three-dimensional imaging of localized surface plasmon resonances of metal nanoparticles *Nature* **502** 80–84
- [127] Xu R *et al* 2015 Three-dimensional coordinates of individual atoms in materials revealed by electron tomography *Nat. Mater.* **14** 1099–103
- [128] Miao J, Ercius P and Billinge S J 2016 Atomic electron tomography: 3D structures without crystals *Science* **353** aaf2157
- [129] Alania M, Altantzis T, De Backer A, Lobato I, Bals S and Van Aert S 2017 Depth sectioning combined with atom-counting in HAADF STEM to retrieve the 3D atomic structure *Ultramicroscopy* **177** 36–42
- [130] Allen J E *et al* 2008 High-resolution detection of Au catalyst atoms in Si nanowires *Nat. Nanotechnol.* **3** 168–73
- [131] Bonnell D A, Kalinin S V, Kholkin A L and Gruverman A 2009 Piezoresponse force microscopy: a window into electromechanical behavior at the nanoscale *MRS Bull.* **34** 648–57
- [132] Moler K A 2017 Imaging quantum materials *Nat. Mater.* **16** 1049–52
- [133] Khajetoorians A A, Wegner D, Otte A F and Swart I 2019 Creating designer quantum states of matter atom-by-atom *Nat. Rev. Phys.* **1** 703–15
- [134] Ziatdinov M, Fujii S, Kiguchi M, Enoki T, Jesse S and Kalinin S V 2016 Data mining graphene: correlative analysis of structure and electronic degrees of freedom in graphenic monolayers with defects *Nanotechnology* **27** 495703
- [135] Susi T, Meyer J C and Kotakoski J 2017 Manipulating low-dimensional materials down to the level of single atoms with electron irradiation *Ultramicroscopy* **180** 163–72
- [136] Gruverman A and Kholkin A 2006 Nanoscale ferroelectrics: processing, characterization and future trends *Rep. Prog. Phys.* **69** 2443–74

- [137] Roccapiore K M, Ziatdinov M, Cho S H, Hachtel J A and Kalinin S V 2021 Predictability of localized plasmonic responses in nanoparticle assemblies *Small* **n/a** 2100181
- [138] Kalinin S V, Kelley K, Vasudevan R K and Ziatdinov M 2021 Toward decoding the relationship between domain structure and functionality in ferroelectrics via hidden latent variables *ACS Appl. Mater. Interfaces* **13** 1693–703
- [139] Yaman M Y, Kalinin S V, Guye K N, Ginger D and Ziatdinov M 2022 Learning and predicting photonic responses of plasmonic nanoparticle assemblies via dual variational autoencoders (arXiv:2208.03861)
- [140] Ziatdinov M, Yaman M Y, Liu Y, Ginger D and Kalinin S V 2021 Semi-supervised learning of images with strong rotational disorder: assembling nanoparticle libraries (arXiv:2105.11475)
- [141] Karniadakis G E, Kevrekidis I G, Lu L, Perdikaris P, Wang S and Yang L 2021 Physics-informed machine learning *Nat. Rev. Phys.* **3** 422–40
- [142] Bepler T, Zhong E, Kelley K, Brignole E and Berger B 2019 Explicitly disentangling image content from translation and rotation with spatial-VAE *Advances in Neural Information Processing Systems* pp 15409–19
- [143] Kalinin S V, Oxley M P, Valletti M, Zhang J, Hermann R P, Zheng H, Zhang W, Eres G, Vasudevan R K and Ziatdinov M 2021 Deep Bayesian local crystallography *npj Comput. Mater.* **7** 181
- [144] Kalinin S V, Steffes J J, Liu Y, Huey B D and Ziatdinov M 2021 Disentangling ferroelectric domain wall geometries and pathways in dynamic piezoresponse force microscopy via unsupervised machine learning *Nanotechnology* **33** 055707
- [145] Tagantsev A K, Cross L E and Fousek J 2010 *Domains in Ferroic Crystals and Thin Films* (Berlin: Springer)
- [146] Morozovska A N, Eliseev E A, Morozovsky N V and Kalinin S V 2017 Ferroionic states in ferroelectric thin films *Phys. Rev. B* **95** 195413
- [147] Highland M J, Fister T T, Fong D D, Fuoss P H, Thompson C, Eastman J A, Streiffer S K and Stephenson G B 2011 Equilibrium polarization of ultrathin PbTiO₃ with surface compensation controlled by oxygen partial pressure *Phys. Rev. Lett.* **107** 187602
- [148] Stephenson G B and Highland M J 2011 Equilibrium and stability of polarization in ultrathin ferroelectric films with ionic surface compensation *Phys. Rev. B* **84** 064107
- [149] Dagotto E 2005 Complexity in strongly correlated electronic systems *Science* **309** 257–62
- [150] Lemos P, Jeffrey N, Cranmer M, Ho S and Battaglia P 2022 Rediscovering orbital mechanics with machine learning (arXiv:2202.02306)
- [151] Udrescu S-M and Tegmark M 2021 Symbolic regression: discovering physical laws from distorted video *Phys. Rev. E* **103** 043307
- [152] Roccapiore K M, Boebinger M G, Dyck O, Ghosh A, Unocic R R, Kalinin S V and Ziatdinov M 2022 Probing electron beam induced transformations on a single-defect level via automated scanning transmission electron microscopy *ACS Nano* **16** 17116–27
- [153] Jia C-L, Mi S-B, Urban K, Vrejoiu I, Alexe M and Hesse D 2008 Atomic-scale study of electric dipoles near charged and uncharged domain walls in ferroelectric films *Nat. Mater.* **7** 57–61
- [154] Riccò S et al 2018 *In situ* strain tuning of the metal-insulator-transition of Ca₂RuO₄ in angle-resolved photoemission experiments *Nat. Commun.* **9** 4535
- [155] Ghosh A, Sumpter B G, Dyck O, Kalinin S V and Ziatdinov M 2021 Ensemble learning-iterative training machine learning for uncertainty quantification and automated experiment in atom-resolved microscopy *npj Comput. Mater.* **7** 1–8
- [156] Kotakoski J, Meyer J C, Kurasch S, Santos-Cottin D, Kaiser U and Krasheninnikov A V 2011 Stone-Wales-type transformations in carbon nanostructures driven by electron irradiation *Phys. Rev. B* **83** 245420
- [157] Liu Y, Proksch R, Wong C Y, Ziatdinov M and Kalinin S V 2021 Disentangling ferroelectric wall dynamics and identification of pinning mechanisms via deep learning *Adv. Mater.* **33** 2103680
- [158] Waser R and Aono M 2007 Nanoionics-based resistive switching memories *Nat. Mater.* **6** 833–40
- [159] Waser R 2012 Nanoelectronics and information technology nanoelectronics and information technology
- [160] Damjanovic D 1998 Ferroelectric, dielectric and piezoelectric properties of ferroelectric thin films and ceramics *Rep. Prog. Phys.* **61** 1267–324
- [161] Setter N et al 2006 Ferroelectric thin films: review of materials, properties, and applications *J. Appl. Phys.* **100** 051606
- [162] Damjanovic D and Demartin M 1997 Contribution of the irreversible displacement of domain walls to the piezoelectric effect in barium titanate and lead zirconate titanate ceramics *J. Phys.: Condens. Matter* **9** 4943
- [163] Seshadri S B, Prewitt A D, Studer A J, Damjanovic D and Jones J L 2013 An *in situ* diffraction study of domain wall motion contributions to the frequency dispersion of the piezoelectric coefficient in lead zirconate titanate *Appl. Phys. Lett.* **102** 042911
- [164] Tutuncu G, Li B, Bowman K and Jones J L 2014 Domain wall motion and electromechanical strain in lead-free piezoelectrics: insight from the model system (1-x) Ba (Zr_{0.2}Ti_{0.8}) O_{3-x} (Ba_{0.7}Ca_{0.3}) TiO₃ using *in situ* high-energy x-ray diffraction during application of electric fields *J. Appl. Phys.* **115** 144104
- [165] Glaum J, Granzow T and Rödel J 2010 Evaluation of domain wall motion in bipolar fatigued lead-zirconate-titanate: a study on reversible and irreversible contributions *J. Appl. Phys.* **107** 104119
- [166] Maksymovych P, Seidel J, Chu Y H, Wu P, Baddorf A P, Chen L-Q, Kalinin S V and Ramesh R 2011 Dynamic conductivity of ferroelectric domain walls in BiFeO₃ *Nano Lett.* **11** 1906–12
- [167] Rojac T et al 2017 Domain-wall conduction in ferroelectric BiFeO₃ controlled by accumulation of charged defects *Nat. Mater.* **16** 322–7
- [168] Kalinin S V, Steffes J J, Liu Y T, Huey B D and Ziatdinov M 2022 Disentangling ferroelectric domain wall geometries and pathways in dynamic piezoresponse force microscopy via unsupervised machine learning *Nanotechnology* **33** 11
- [169] Kalinin S V, Rama Vasudevan A G and Ziatdinov M 2021 From atomically resolved imaging to generative and causal models *Nat. Phys.* **18** 1152–60
- [170] Ayana Ghosh G P, Trujillo D P, Shaikh M and Ghosh S 2022 Insights into cation ordering of double perovskite oxides from machine learning and causal relations *Chem. Mater.* **34** 7563–78
- [171] Maxim Ziatdinov A G and Kalinin S V 2021 Physics makes the difference: Bayesian optimization and active learning via augmented Gaussian process *Mach. Learn.: Sci. Technol.* **3** 015022
- [172] Hourahine B et al 2020 DFTB+, a software package for efficient approximate density functional theory based atomistic simulations *J. Chem. Phys.* **152** 124101
- [173] Ayana Ghosh M Z, Dyck O, Sumpter B and Kalinin S V 2021 Bridging microscopy with molecular dynamics and quantum simulations: an AtomAI based pipeline *npj Comput. Mater.* **8** 11
- [174] Mody C C M 2011 *Instrumental Community* (Cambridge, MA: The MIT Press)
- [175] Ziatdinov M, Maksov A and Kalinin S V 2017 Learning surface molecular structures via machine vision *npj Comput. Mater.* **3** 31

- [176] Ziatdinov M, Dyck O, Li X, Sumpter B G, Jesse S, Vasudevan R K and Kalinin S V 2019 Building and exploring libraries of atomic defects in graphene: scanning transmission electron and scanning tunneling microscopy study *Sci. Adv.* **5** eaaw8989
- [177] Maksov A, Dyck O, Wang K, Xiao K, Geohegan D B, Sumpter B G, Vasudevan R K, Jesse S, Kalinin S V and Ziatdinov M 2019 Deep learning analysis of defect and phase evolution during electron beam-induced transformations in WS₂ *npj Comput. Mater.* **5** 1–8
- [178] Gordon O, D'Hondt P, Knijff L, Freeney S E, Junqueira F, Moriarty P and Swart I 2019 Scanning tunneling state recognition with multi-class neural network ensembles *Rev. Sci. Instrum.* **90** 103704
- [179] Farley S et al 2020 Improving the segmentation of scanning probe microscope images using convolutional neural networks *Mach. Learn.: Sci. Technol.* **2** 015015
- [180] Nguyen T N M, Guo Y, Qin S, Frew K S, Xu R and Agar J C 2021 Symmetry-aware recursive image similarity exploration for materials microscopy *npj Comput. Mater.* **7** 14
- [181] Ziatdinov M, Nelson C, Vasudevan R K, Chen D and Kalinin S V 2019 Building ferroelectric from the bottom up: the machine learning analysis of the atomic-scale ferroelectric distortions *Appl. Phys. Lett.* **115** 052902
- [182] Kalinin S V, Dyck O, Jesse S and Ziatdinov M 2021 Exploring order parameters and dynamic processes in disordered systems via variational autoencoders *Sci. Adv.* **7** eabd5084
- [183] Oxley M P, Ziatdinov M, Dyck O, Lupini A R, Vasudevan R and Kalinin S V 2021 Probing atomic-scale symmetry breaking by rotationally invariant machine learning of multidimensional electron scattering *npj Comput. Mater.* **7** 65
- [184] Liu Y, Kelley K P, Vasudevan R K, Funakubo H, Ziatdinov M A and Kalinin S V 2022 Experimental discovery of structure–property relationships in ferroelectric materials via active learning *Nat. Mach. Intell.* **4** 341–50
- [185] Roccapriore K M, Dyck O, Oxley M P, Ziatdinov M and Kalinin S V 2022 Automated experiment in 4D-STEM: exploring emergent physics and structural behaviors *ACS Nano* **16** 7605–14
- [186] Wilson A G, Hu Z, Salakhutdinov R and Xing E P 2016 Artificial intelligence and statistics *PMLR* pp 370–8
- [187] Kalinin S V, Ziatdinov M, Hinkle J, Jesse S, Ghosh A, Kelley K P, Lupini A R, Sumpter B G and Vasudevan R K 2021 Automated and autonomous experiment in electron and scanning probe microscopy (arXiv:2103.12165)
- [188] Cichos F, Gustavsson K, Mehlig B and Volpe G 2020 Machine learning for active matter *Nat. Mach. Intell.* **2** 94–103
- [189] Krull A, Hirsch P, Rother C, Schiffrin A and Krull C 2020 Artificial-intelligence-driven scanning probe microscopy *Commun. Phys.* **3** 54
- [190] Ede J M 2020 Warwick electron microscopy datasets *Mach. Learn.: Sci. Technol.* **1** 045003
- [191] Wei J, Francis C, Morgan D, Schmidt K J, Scourtas A, Foster I, Blaiszik B and Voyles P M 2022 Infrastructure for analysis of large microscopy and microanalysis data sets *Microsc. Microanal.* **28** 3094–6
- [192] Borisevich A et al 2010 Mapping octahedral tilts and polarization across a domain wall in BiFeO₃ from Z-contrast scanning transmission electron microscopy image atomic column shape analysis *ACS Nano* **4** 6071–9
- [193] He Q, Ishikawa R, Lupini A R, Qiao L, Moon E J, Ovchinnikov O, May S J, Biegalski M D and Borisevich A Y 2015 Towards 3D mapping of BO₆ octahedron rotations at perovskite heterointerfaces, unit cell by unit cell *ACS Nano* **9** 8412–9
- [194] Steiner S et al 2019 Organic synthesis in a modular robotic system driven by a chemical programming language *Science* **363** eaav2211
- [195] Krenn M, Häse F, Nigam A, Friederich P and Aspuru-Guzik A 2020 Self-referencing embedded strings (SELFIES): a 100% robust molecular string representation *Mach. Learn.: Sci. Technol.* **1** 045024
- [196] Stocker H, Perna L, Weigl K, Möllers T, Schöttker B, Thomsen H, Holleczeck B, Rujescu D and Brenner H 2021 Prediction of clinical diagnosis of Alzheimer's disease, vascular, mixed, and all-cause dementia by a polygenic risk score and APOE status in a community-based cohort prospectively followed over 17 years *Mol. Psychiatry* **26** 5812–22
- [197] Coley C W, Green W H and Jensen K F 2018 Machine learning in computer-aided synthesis planning *Acc. Chem. Res.* **51** 1281–9
- [198] Wang W R, Jiang X, Tian S, Liu P, Dang D, Su Y, Lookman T and Xie J 2022 Automated pipeline for superalloy data by text mining *npj Comput. Mater.* **8** 12
- [199] Olivetti E A, Cole J M, Kim E, Kononova O, Ceder G, Han T Y-J and Hiszpanski A M 2020 Data-driven materials research enabled by natural language processing and information extraction *Appl. Phys. Rev.* **7** 19
- [200] Tshitoyan V, Dagdelen J, Weston L, Dunn A, Rong Z, Kononova O, Persson K A, Ceder G and Jain A 2019 Unsupervised word embeddings capture latent knowledge from materials science literature *Nature* **571** 95
- [201] Huo H Y, Rong Z, Kononova O, Sun W, Botari T, He T, Tshitoyan V and Ceder G 2019 Semi-supervised machine-learning classification of materials synthesis procedures *npj Comput. Mater.* **5** 7
- [202] Epps R W, Felton K C, Coley C W and Abolhasani M 2017 Automated microfluidic platform for systematic studies of colloidal perovskite nanocrystals: towards continuous nano-manufacturing *Lab Chip* **17** 4040–7
- [203] Symes M D, Kitson P J, Yan J, Richmond C J, Cooper G J T, Bowman R W, Vilbrandt T and Cronin L 2012 Integrated 3D-printed reaction ware for chemical synthesis and analysis *Nat. Chem.* **4** 349–54
- [204] Coley C W 2021 Defining and exploring chemical spaces *Trends Chem.* **3** 133–45
- [205] Wu T L and Tegmark M 2019 Toward an artificial intelligence physicist for unsupervised learning *Phys. Rev. E* **100** 19
- [206] Udrescu S-M, Tegmark M and Feynman: A I 2020 A physics-inspired method for symbolic regression *Sci. Adv.* **6** eaay2631
- [207] Udrescu S-M et al 2020 AI Feynman 2.0: pareto-optimal symbolic regression exploiting graph modularity *Advances in Neural Information Processing Systems* vol 33 pp 4860–71
- [208] Champion K, Lusch B, Kutz J N and Brunton S L 2019 Data-driven discovery of coordinates and governing equations *Proc. Natl Acad. Sci.* **116** 22445–51
- [209] Brunton S L, Proctor J L and Kutz J N 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems *Proc. Natl Acad. Sci.* **113** 3932–7
- [210] Valletti S M P et al 2021 Bayesian learning of adatom interactions from atomically resolved imaging data *ACS nano* **15** 9649–57
- [211] Schölkopf B, Locatello F, Bauer S, Ke N R, Kalchbrenner N, Goyal A and Bengio Y 2021 Toward causal representation learning *Proc. IEEE* **109** 612–34
- [212] Kingma D P and Welling M 2019 An introduction to variational autoencoders (arXiv:1906.02691)
- [213] Keen D A and Goodwin A L 2015 The crystallography of correlated disorder *Nature* **521** 303–9
- [214] Overy A R, Cairns A B, Cliffe M J, Simonov A, Tucker M G and Goodwin A L 2016 Design of crystal-like aperiodic solids with selective disorder-phonon coupling *Nat. Commun.* **7** 10445
- [215] Kunitski M et al 2019 Kinetic control of tunable multi-state switching in ferroelectric thin films *Nat. Commun.* **10** 1–10
- [216] Korolkov V V, Summerfield A, Murphy A, Amabilino D B, Watanabe K, Taniguchi T and Beton P H 2019 Ultra-high resolution imaging of thin films and single strands of polythiophene using atomic force microscopy *Nat. Commun.* **10** 1–8

- [217] Vasudevan R K et al 2021 Investigating phase transitions from local crystallographic analysis based on statistical learning of atomic environments in 2D MoS₂-ReS₂ *Appl. Phys. Rev.* **8** 011409
- [218] Butler K T, Davies D W, Cartwright H, Isayev O and Walsh A 2018 Machine learning for molecular and materials science *Nature* **559** 547–55
- [219] Muratov E N et al 2020 QSAR without borders *Chem. Soc. Rev.* **49** 3525–64
- [220] Cadeddu A, Wylie E K, Jurczak J, Wampler-Doty M and Grzybowski B A 2014 Organic chemistry as a language and the implications of chemical linguistics for structural and retrosynthetic analyses *Angew. Chem., Int. Ed.* **53** 8108–12
- [221] Ebbesen T W, Lezec H J, Ghaemi H, Thio T and Wolff P A 1998 Extraordinary optical transmission through sub-wavelength hole arrays *Nature* **391** 667–9
- [222] Huang X and El-Sayed M A 2011 Plasmonic photo-thermal therapy (PPTT) *Alexandria J. Med.* **47** 1–9
- [223] Tran T T, Elbadawi C, Totonjian D, Lobo C J, Grosso G, Moon H, Englund D R, Ford M J, Aharonovich I and Toth M 2016 Robust multicolor single photon emission from point defects in hexagonal boron nitride *ACS Nano* **10** 7331–8
- [224] Kociak M and Stéphan O 2014 Mapping plasmons at the nanometer scale in an electron microscope *Chem. Soc. Rev.* **43** 3865–83
- [225] Way M 2017 What I cannot create, I do not understand *J. Cell. Sci.* **130** 2941–2
- [226] Jesse S, Borisevich A Y, Fowlkes J D, Lupini A R, Rack P D, Unocic R R, Sumpter B G, Kalinin S V, Belianinov A and Ovchinnikova O S 2016 Directing matter: toward atomic-scale 3D nanofabrication *ACS Nano* **10** 5600–18
- [227] Gonzalez-Martinez I G, Bachmatiuk A, Bezugly V, Kunstmann J, Gemming T, Liu Z, Cuniberti G and Rummeli M H 2016 Electron-beam induced synthesis of nanostructures: a review *Nanoscale* **8** 11340–62
- [228] Fuechsle M, Miwa J A, Mahapatra S, Ryu H, Lee S, Warschkow O, Hollenberg L C L, Klimeck G and Simmons M Y 2012 A single-atom transistor *Nat. Nanotechnol.* **7** 242–6
- [229] Feynman R P 1960 There's plenty of room at the bottom *Caltech Eng. Sci.* **23** 22–36
- [230] Kalinin S V, Jesse S and Lupini A R 2022 A quantum lab in a beam *Phys. Today* **75** 30–36
- [231] Tripathi M, Mittelberger A, Pike N A, Mangler C, Meyer J C, Verstraete M J, Kotakoski J and Susi T 2018 Electron-beam manipulation of silicon dopants in graphene *Nano Lett.* **18** 5319–23